ESC-TR-95-003

Technical Report
1020

A Satellite Anomaly Detection System that Uses Signatures from the Space-Surveillance Network of Sensors

R.C. Raup
J.S. Stuart
J.W. Curtis

7 September 1995

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

Prepared for the Department of the Air Force under Contract F19628-95-C-0002.

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the Department of the Air Force under Contract F19628-95-C-0002.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Gary Tutungian

Gary Tutungian
Administrative Contracting Officer
Contracted Support Management

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

A SATELLITE ANOMALY DETECTION SYSTEM THAT USES SIGNATURES FROM THE SPACE-SURVEILLANCE NETWORK OF SENSORS

R.C. RAUP
J.S. STUART
J.W. CURTIS
Group 91

Accesion For		
NTIS CRA&I		☒
DTIC TAB		☐
Unannounced		☐
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

TECHNICAL REPORT 1020

7 SEPTEMBER 1995

Approved for public release; distribution is unlimited.

LEXINGTON MASSACHUSETTS

ABSTRACT

Lincoln Laboratory has developed a capability to assess and monitor the status of high-earth-orbit (HEO) satellites using simple radar-derived satellite target signatures. The capability takes the form of a computer workstation that processes, stores, and retrieves HEO satellite signatures and associated data products. The important sensor phenomena, signal and data processing algorithms, software architecture, and concept of operations are all described in this report.

The system was designed and tested using L-band radar data, but will accept other narrowband radar data and photometric data as well. Signal and data processing algorithms are based on statistical modeling and methods. The amount of cross-section fluctuation in each signature is measured as a function of various independent observation variables. The fluctuation measurements are first used to characterize the normal operation of the satellite and then to provide a basis for automatic detection of anomalies. Thus changes in satellite status that manifest changes in the apparent stability of the satellite are detected by the system. The algorithms are embedded in a network-transparent UNIX software architecture with an MIT X Windows and Motif interface. The software for the system can be used on any UNIX platform with little or no modification.

The basic facility may consist of one or more contributing radars or photometric sensors, a data processing center, and one or more analysis centers. These facility elements need not reside at the same location because the software may be distributed among several computing platforms at diverse geographical locations. Signal processing, data processing, database management, and anomaly detection are automated, while database retrieval has both an automated mode for sensor control-room displays and an interactive mode for data analysis.

ACKNOWLEDGMENTS

The contributions of Stephen M. Hunt to the early stages of this project are appreciated, particularly his contribution to the prototype of the analysis interface.

TABLE OF CONTENTS

TABLE OF CONTENTS
(Continued)

LIST OF ILLUSTRATIONS

LIST OF TABLES

1. INTRODUCTION

Space surveillance is a means of gaining knowledge of the location, function, and status of the man-made satellites in space. The number of satellites continues to increase even as the resources available for space surveillance, both the men and machines, continues to decline. It is clear that the techniques currently used in space surveillance are too expensive to survive all future budget cuts. Somehow the important products of space surveillance must be produced more efficiently in the future.

The U.S. utilizes a worldwide network of sensors, the space-surveillance network (SSN), to routinely track many thousands of earth-orbiting satellites. These sensors determine the positions of the satellites with sufficient quality to maintain a catalog describing the orbit of each satellite. The orbit is a mathematical description of the translation (as opposed to rotation) of the satellite that is determined from measurements of the satellite's position over time. The orbit can be used to predict the position of the satellite many days into the future or the past. Predicting a satellite's position is the main reason for determining the orbit of a satellite. The accuracy of the prediction decreases as the prediction is extended farther into the future or past. The orbits can be used to predict the location of each satellite for days or weeks into the future, thus helping to keep track of a growing inventory of orbiting payloads and debris.

The determination of orbits is an important space-surveillance product. Orbit determination will continue to be important in the future. Satellites can be designed to monitor and report their own positions or otherwise actively support the determination of their orbits by ground-based equipment. But even if all future satellites carry such electronic systems, the orbits of inactive, damaged, and uncooperative satellites must be monitored by space-surveillance techniques. This is the main justification for conducting space surveillance.

Some sensors within the SSN, such as radars and optical telescopes with photometers, record the power reflected from the satellite (either microwave or optical power) over time as the satellite is tracked to maintain the catalog of orbits. These power-time series, called signatures, are not routinely used by the SSN. Given a need to maintain a catalog of orbits, it is interesting to consider if the signatures, a by-product of tracking the satellite for orbital catalog maintenance, can be used to determine if a satellite is functioning normally. This report describes a system for satellite assessment and monitoring that uses the signatures already available from the SSN.

The report continues with an introductory Section 2 that describes the features of space-surveillance sensors and measurements that are needed in order to understand the remainder of the report. The report details the use of noncoherent narrowband radar data, but certain aspects of optical sensors are also covered in Section 2. This is because photometric data is similar to noncoherent narrowband radar data in many ways. The remainder of the report addresses only the radar case with the hope that the reader will be able to understand the implications to optical sensors after reading this section. In Section 3 some of the relevant characteristics of satellite observations that are made by narrowband radars are discussed, laying the groundwork for Section 4, where satellite characterization is introduced. Section 4 describes the basic statistical concepts behind the techniques used to characterize high-altitude satellites by the sensor measurements. The intuitive details of signal and data processing are covered in Section 5. In Section 6 some of the engineering considerations used to design the workstation implementation of the system are considered. Although the system operates autonomously, a human analyst provided with a graphical user interface (GUI) is expected to review the anomaly detection messages produced by the system and make the final decisions regarding the status of satellites. This concept of operation is described in Section 7. A statistical summary of the database in a typical operating system is presented in Section 8.

2. SPACE-SURVEILLANCE SENSORS AND MEASUREMENTS

In this section the focus is primarily on monostatic microwave radars and passive optical sensors that use visible light (not infrared radiation). There are two reasons why these generic sensors are investigated. First, they are the primary space-surveillance sensors. Radars originally built to detect strategic bombers or ballistic missiles have evolved into the space-surveillance task, while the optical sensors were primarily designed for space surveillance. The two types of sensors offer complementary capabilities that are described in this section. Second, most of the surveillance techniques emphasized in this report apply to both radar data and optical-sensor data. Both types of data can be obtained by measuring the power reflected from the satellite as a function of time, that is, the satellite signature. In the case of the radar, the energy is supplied by a microwave transmitter, and in the case of the optical sensor the energy is reflected sunlight.

This section will compare the basic measurement capabilities of the radars and the optical sensors. The basic measurement capabilities associated with the two types of sensors are similar in some respects and very different in others. These similarities and differences are pointed out. The main goal of Section 2 is to provide the self-contained overview of the sensor characteristics that is required to understand the remainder of the report. The following sections will emphasize the radar, but the discussion here should make the analogies between radars and optical sensors fairly clear. After reading this section, the careful reader will be able to imagine how the other sections about radars can be applied to optical sensors.

2.1 Space-Surveillance Sensors

Space surveillance now is accomplished primarily by two types of ground-based sensors—microwave radars and passive optical sensors. A microwave radar emits microwave energy that is reflected from the satellite and measured as it returns to the radar. Because it illuminates the satellite with its own energy, a radar is called *active*. A passive optical sensor uses a telescope-like device to record the reflected sunlight or infrared energy emitted by the satellite. The sensor does not generate the light or infrared energy, and so it is called *passive*.

Microwave radiation, visible light, and infrared radiation are all forms of electromagnetic radiation. Figure 1 shows the frequencies (and equivalently the wavelengths) used to characterize electromagnetic radiation, as well as the common names of electromagnetic radiation bands. Most passive optical systems use either infrared or visible light, but most available space-surveillance sensors operate in the visible band.

The component of the microwave radar that emits the microwave energy is called the *transmitter*. If the passive optical sensor measures the sunlight reflected from the satellite, the sun can be thought of as a transmitter too, although one that is obviously not under the control of the sensor. The sensor component that measures the power from the satellite is called a *receiver*. When the transmitter and the receiver are located at essentially the same spot, the sensor is called

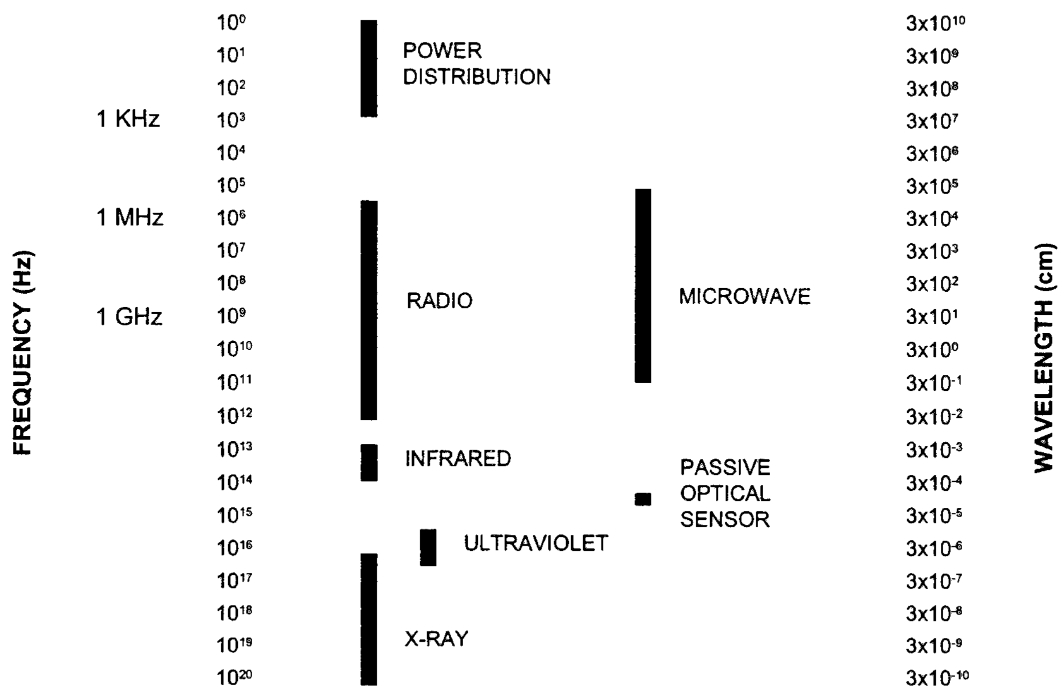

Figure 1. Two common ground-based satellite surveillance sensors, the microwave radar and the passive optical sensor using visible light, operate at frequencies that differ by five orders of magnitude.

monostatic, otherwise it is *bistatic*. Most space-surveillance radars are monostatic. The satellite, the properties of which are being measured by a sensor, is called the sensor's *target*.

2.2 Sensor Measurements

Radars and optical sensors measure both the satellite's cross section and certain parameters related to the satellite's position. These two types of measurements are described in this section.

2.2.1 Satellite Cross Section

Simply put, the concept of *cross section* is mathematically a factor (in units of area, such as square meters) used to balance the equations describing the conservation of power that flows from the satellite to the sensor's receiver. A formal discussion can be found in Section 1.4 of Bhattacharyya and Sengupta [1]. The cross section associated with the microwave radar is called the satellite's *radar cross section*, while the cross section associated with the optical sensor is called the satellite's *optical cross section*.

Under idealized conditions, the amount of power radiated by the satellite is proportional to the *physical cross section* obtained by projecting the portion of satellite illuminated by the transmitter onto a plane normal to the line of sight between the satellite and the receiver. This physical cross section is similar to a shadow, and is undoubtedly the source of the expressions "radar cross section" and "optical cross section." But a radar cross section or optical cross section does not generally have the same value as the physical cross section. In fact, all three notions of cross section generally take on different values for the same given satellite. Furthermore, the cross section measured with a monostatic sensor will be different from the cross section measured with a bistatic sensor.

Thus, sensor-measured cross sections provide a crude indication of the physical size of the satellite at best. The cross sections depend on the aspects from which the satellite is illuminated and viewed. As these angles change a small amount, the sensor-measured cross sections can change very rapidly. The sensitivity of the cross section to small changes in angle is perhaps most evident for radars, which generally illuminate the satellite with *coherent* radiation.

Here are the consequences of radiation coherency that are most important in this report. Natural light is a so-called *noncoherent* source of illumination. This means that the power seen by the optical sensor's receiver is a simple addition of the powers returned by each little patch on the satellite surface. The nature of coherent radiation implies that this simple power additivity property will not hold. The energy reradiated by the radar-illuminated satellite may result in very little power (or maybe even none) at the receiver when the satellite is viewed from certain aspects. There are now entire books on this subject, such as Bhattacharyya and Sengupta [1].

The radar or optical cross section is determined from the amount of power that enters the sensor's receiver. Thus estimates of the satellite cross section are intimately related to the satellite's signature, which is a power-time series measured by the sensor. The cross section is estimated by suitably scaling the power received by the sensor. The signature of the satellite is often presented in this scaled form as the time-varying satellite cross section.

2.2.2 Satellite Position

The angular position of the satellite within the field of view can be determined with both radars and optical sensors. The *field of view* (FOV) is the volume of space from which radiation enters the sensor's receiver. Because the field of view is often conical in shape, the field of view can be expressed as the angle that determines the cone.

In the optical sensor the position of the satellite on the sensor's focal plane is used to measure the angular difference between the direction to the satellite and the instrument's boresight (which is represented by a particular point on the focal plane) or the angle between the direction to the satellite and the direction to a known star in the field of view. The *boresight* is an imaginary line from the sensor, usually centered in its field of view. When the field of view is conical, the boresight is usually the axis of the cone. The *focal plane* of an optical sensor is the plane where the receiver projects the magnified image of its field of view, similar to the film plane in a camera.

For a radar there is no detectable object in the field of view that is analogous to a star with an a priori known position. Generally only the satellite is detectable in the field of view. But different phase shifts (very small time delays) in the coherent radiation received by the radar will occur over the physical extent of the antenna. The value of the phase shifts depends on the angular difference between the satellite's position and the radar's boresight. The phase shifts can be measured and used to estimate the angle between the direction to the satellite and the boresight (see, for example, Chapter 14 of Levanon [2] or pages 207–238 of Carpentier [3]). The boresight orientation must be known to get the absolute direction to the satellite.

Because the radar controls its transmitter, it can measure the propagation time of the energy as it traverses the space to the satellite and back, yielding a direct estimate of the range to the satellite. The propagation velocity of the microwave radiation is known fairly accurately in both the vacuum of space and within the earth's atmosphere. A satellite range estimate cannot be directly produced by an optical sensor, which cannot time the transit of the sunlight reflected by the satellite.

Over time, the measurements of the position of the satellite are used to determine the satellite's orbit. This is the principal utility of the position measurements. Once the orbit is determined it is very easy to locate the satellite again because its position can be predicted accurately for some time.

2.2.3 Other Measurements

Certain rarely used analyses extract additional information about the satellite. The polarization of the energy from the satellite can be measured by some radars and optical sensors. The *polarization* of the energy provides information about the orientation of the electric and magnetic vectors that determine the nature of the electromagnetic radiation (see, for example, Section 1.4.2 of Born and Wolf [4]). Also, because the optical illumination is noncoherent and has a large bandwidth, the optical sensor can make a measure of the color of the satellite. Both polarization and color become additional tools, along with cross section, to characterize the satellite. These other measurements will not be pursued in this report.

2.3 Comparison of Microwave Radars and Passive Optical Sensors

Fundamental differences between radars and optical sensors follow from the different frequencies of satellite illumination, the different degree of coherence in the illuminating energy, and the illumination and viewing geometries.

One obvious difference involves cost, weight, and power requirements. The antenna and transmitter components of the radar are bulky and expensive compared to the telescopes used by the optical sensors, making optical sensors generally cheaper, more compact, and more transportable. The radar transmitter is power-hungry, consuming several megawatts of power.

The radar transmitter and receiver are generally collocated (monostatic). The sun and the optical sensor are obviously not collocated (bistatic). This complicates reconciliation of the measurements of the radar and optical cross sections of the satellite because different parts of the satellite are illuminated depending on which type of sensor is used.

Finally the radar radiation is coherent so that the power received by the radar cannot be modeled as a sum of the incremental powers originating from each tiny patch of illuminated target. This may make the amount of power received by the radar highly dependent on the aspect angle from which the target is viewed and further complicates the reconciliation of radar, optical, and physical target cross sections.

3. RADAR OBSERVATIONS OF SATELLITES

The principal measurements obtained from a space-surveillance sensor are the satellite cross section and the satellite position as a function of time. These constitute the radar observations of satellites that are relevant to this report. By Section 4 the emphasis will be on the cross section measurements.

As discussed in Section 2.2.1, the cross section measurement is not really any physical cross section of the satellite, but it is definitely related to the physical structure of the satellite. As such, studying the cross section measurements of a satellite is best done with knowledge of the viewing aspect at which the cross section measurement was made. This requires knowledge of the orientation of the satellite.

The measurements of the position of the satellite are essentially the same as a knowledge of the satellite's orbit. It is not the same as knowledge of the satellite's orientation. The more difficult problem of determining orientation directly is usually addressed with data from wideband radars. In some cases orientation can be determined by extensive modeling in conjunction with narrowband radar or optical-sensor data, but this is not routinely done.

It is difficult at best to deduce the orientation of the satellite from narrowband radar observations. In this section some of the operational features of satellites are discussed to convince the reader that the orientation of the satellite, although not directly deducible from the radar observations in most cases, is related to certain independent variables, the values of which can be calculated by the observer. This happens because the relation is part of the design and normal operation of the satellite. It leads to some imprecise constraints on the possible orientations of satellites that are utilized in the satellite characterization techniques of Section 4. With this objective in mind, Section 3.1 discusses some of the relationships between satellite structure, satellite operation, and the observed cross section and position of the satellite. Section 3.2 discusses some of the relevant issues related to the quality of the radar observations of satellites.

3.1 Satellite Structure and Operation

3.1.1 Orbital Classes

Satellite orbits are often divided into at least two regimes: the *low-earth orbit* (LEO) and the *high-earth orbit* (HEO). The lower the altitude of the satellite, the shorter its orbital period. A satellite in LEO is visible many times each day and its entire passage from horizon to horizon can be observed in a few minutes. The LEO always has a small value of eccentricity (is nearly circular), otherwise increased drag at perigee will significantly reduce the orbital lifetime, and the satellite will reenter the earth's atmosphere, usually within weeks. The eccentricity describes the ellipticity, or deviation from circularity, of the orbit. At 200 km a satellite with a circular orbit will have a lifetime substantially less than the satellite with a high-eccentricity orbit and a 200-km perigee. The *perigee* is the point on the orbital ellipse closest to the primary focus. The *apogee* is the point on the orbital ellipse farthest from the primary focus. Although both satellites pass within 200 km of the ground, the satellite with the high-eccentricity orbit spends significant time at higher altitudes where drag is less, extending its orbital lifetime. An elliptical orbit is used when both long lifetime and some close-earth approaches are required for the satellite's mission.

In a high-eccentricity orbit the true anomaly of the satellite becomes an important independent variable for studying the radar observations of satellites. The *true anomaly* is the angle measured from the primary focus of the orbital ellipse (roughly the earth's center) between the direction to the satellite and the direction to the perigee. Some satellites are designed to perform certain operations as a function of the true anomaly: one operation performed at apogee, for example, another operation performed at perigee.

For true LEO satellites, which have nearly circular orbits at low altitude, the true anomaly variable is not a very interesting independent variable to use for characterization of the radar observations of satellites. It is undefined for an exactly circular orbit, and estimates of the perigee are very noisy for orbits that are only slightly noncircular. Some satellites in LEO keep antennas or other components directed toward the center of the earth. Then the aspect from which the sensor views this satellite component depends on the apparent elevation of the satellite above the horizon as seen from the sensor. For such satellites the elevation becomes an important independent variable for studying the radar observations.

A satellite in HEO is generally visible to a point on the earth for very long periods of time. As a special case it is geosynchronous, that is, for most points on the earth it is either always visible or always below the horizon. This is accomplished by setting the orbital period equal to the earth's rotation time. If its orbital plane does not exactly contain the earth's equator or the orbit is not exactly circular, some apparent motion of a geosynchronous satellite will be observed from the ground, essentially repeating every 24 hours. The HEO can have large eccentricities (be very elliptical) that, in the extreme, keep the perigee height only as high as it needs to be to prevent premature reentry of the satellite.

A third regime is sometimes defined—the mid-earth orbit (MEO). In MEO a satellite's apparent motion from horizon to horizon requires several hours. As a special class, MEOs contain the half-synchronous satellites that to the earth-based observer appear to pass from horizon to horizon twice a day in a pattern repeating exactly every 24 hours. The half-synchronous satellites orbit with a period equal to one-half of the earth's rotation time.

Both the synchronous and half-synchronous satellites trace the same ground path over the surface of the earth every 24 hours. A point on the earth is on the *ground path* of the satellite when at some time in its orbit the satellite appears to be directly overhead. For these satellites, the ground paths of which repeat every 24 hours, the latitude and longitude of the point on the ground path corresponding to the satellites' position become important independent variables for studying radar observations of satellites. These latitude and longitude values are approximately the *geodetic coordinates* of the satellite position. Some satellites are designed to perform certain operations as a function of their geodetic location. Also, because the ground paths repeat every

24 hours, the time of day can be used just as well as the geodetic coordinates, because a particular time of day corresponds to a particular geodetic coordinate.

Table 1 summarizes some of the independent variables, arising from the different orbital classes, that are useful for studying radar observations of satellites. In each orbital class the most useful variables are marked in the appropriate column. The satellite characterization techniques of

TABLE 1

Variables Associated with Different Orbital Classes

	Time of Day	True Anomaly	Satellite Elevation	Geodetic Coordinates
LEO			X	X
Eccentric MEO		X		
Half Synch MEO	X			X
Eccentric HEO		X		
Synchronous HEO	X			X

Section 4 require that the satellite not appear to move very much along its orbit while its signature is collected by the sensor. Hence most of the discussion in the remainder of the report concerns the satellites in HEO and MEO. There are other important variables that arise more directly from the stabilization references that satellite attitude control systems utilize.

3.1.2 Stabilization References

Active satellites must orient sensors, antennas, and solar cells. They must protect certain components from direct sun. So, satellites usually assess and control their orientation. Table 2 lists some of the references commonly used by satellite attitude control systems. Because the satellite is designed to control its orientation or configuration with respect to one or more of these references, the references become a source of valuable independent variables for studying radar observations of satellites. The *configuration* of the satellite describes the relative orientation of separately movable components. These satellite components may include the main body, some antennas, and the solar-panel array.

For example, the angle defined by the two vectors from the satellite to the sun and the satellite to the sensor is often related to the aspect from which the sensor views the solar-panel structure of a solar-powered satellite. This angle, called the *phase angle*, will be important later in the report.

TABLE 2

Attitude Control References

Reference	Implementation	Comments
Sun	Needed to supply solar power to the satellite.	Constrains the position of the solar panels.
Earth	Needed to maintain communications with earth and perform earth-monitoring functions.	Constrains the positions of the antennas and the sensor heads.
Inertial	Needed for precise pointing.	Maintains the satellite orientation in earth shadow and during standby conditions. Constrains the orientation of the main satellite body.

The antennas of the satellite may be kept pointing at the earth. The main body of the satellite may be inertially stabilized while the solar-panel arrays and antennas are positioned via additional degrees-of-motion freedom. Such satellites are said to be able to alter their configuration.

3.2 Radar Calibration

Calibration is the process of determining and maximizing the accuracy of the radar measurements. Calibration must be applied (1) to the position measurements that are used to estimate the satellite's orbit and (2) to the measurements of received power that determine the satellite's signature. This report is concerned with the calibration issues affecting the signature, because most space-surveillance radars can measure satellite positions accurately enough to apply the techniques of Sections 4 and 7.

The power measured in the radar receiver can be related to the satellite's cross section, a quantity fundamentally related to the satellite's composition and structure. Hence the first thought is to estimate the satellite's cross section from the radar measurements and use the estimated time-varying satellite cross section as part of any attempt to discover intrinsic characteristics of the satellite. However this may not be practical for the following reasons.

The calculation of the satellite cross section from the receiver power requires knowledge of several important quantities. Somehow the amount of power incident on the satellite and the amount of power reflected from the satellite need to be determined. Such things as the actual power radiated by the radar, the losses in converting the energy reflected from the satellite into power in the receiver, and so forth become crucial. Because of the practical errors in the knowledge of such quantities, the estimates of the cross section derived from current space-surveillance sensors are often in error by even one hundred percent!

It is difficult to compare the absolute cross-section measurements from two different radars or from the same radar on two different days. Some optimists would claim that errors are smaller than this, but in practice cross-section errors will not be smaller than one hundred percent routinely and consistently. Perhaps it is because the space-surveillance sensors have been primarily used to provide satellite position measurements, and the need has not arisen to provide better signature calibration. At any rate, any data processing scheme that is sensitive to cross-section estimate errors is not likely to function well with the current space-surveillance system.

One does observe, however, that changes in cross section as a function of viewing aspect match well in simultaneous observations of a satellite by two radars or by two observations from the same radar observing the same change in satellite aspect on different days. That is, changes in cross section are less susceptible to calibration problems than the absolute value of cross section.

This is because within a single observation of a satellite, lasting perhaps 10 minutes or less, the dominant error in cross-section estimates is a bias. A *bias* is a constant error that does not change for at least several minutes. Data processing schemes that are not sensitive to bias errors will function better than schemes sensitive to bias errors with the current space-surveillance system. The methods described in Section 4 are insensitive to bias errors.

4. HEO SATELLITE CHARACTERIZATION WITH NARROWBAND RADARS

Narrowband radars cannot provide direct structural details of a satellite. Is it still possible that the narrowband radar signature data could be reliably and automatically applied to the satellite status monitoring problem? Much of this data is available because the SSN radars are used to determine the orbits of satellites. Most of the signature data collected simultaneously with the satellite position measurements used to compute the orbits is currently underutilized or discarded.

Because the narrowband radars cannot be routinely used to directly extract the physical size, shape, and orientation of the satellite, some other approach must be used. A few ideas follow from recognizing that satellites exhibit *nonsmooth* and *nonrandom* behavior as a function of some of the variables discussed in Section 3.1. These ideas will lead to a method of characterizing a satellite in this section. In Section 4 the algorithms that are used to implement the satellite characterization and status monitoring concepts will be developed.

In Section 2.2.1 the radar cross-section measurement was introduced. One of the key ideas was that the cross section was highly dependent on the aspect angle from which the sensor viewed the satellite. The exact viewing aspect is not generally known, but when the satellite is always viewed from the same aspect angle, the cross section is the same unless the satellite changes its configuration. In Section 3.1 a number of independent variables, the values of which are known to the sensor, were introduced. Arguments were presented that loosely linked these variables to constraints on the orientations of the various satellite components that contribute to the satellite cross section. This section combines the two ideas to illustrate a satellite characterization scheme utilizing the independent variables to parameterize the amount of fluctuation in the narrowband radar cross-section measurements of satellites. By using cross-section fluctuation instead of the absolute cross section values, the characterization is less susceptible to the bias errors discussed in Section 3.2.

4.1 Smooth and Nonsmooth Behavior

The external forces that influence the dynamic motion of a satellite, principally gravity and low-altitude drag, vary smoothly over time and space. Because acceleration is proportional to force, the dynamic motion of a satellite influenced only by external forces is also smooth. Its orbit changes slowly with time. Its rotation rate and axis change slowly with time and space.

Figure 2(a) illustrates the smooth change in the geodetic longitude of an inactive geosynchronous satellite. The longitude changes due to natural forces. Figure 2(b) illustrates the nonsmooth variation in geodetic longitude resulting from the periodic thrust applied by the satellite to maintain a desired value of longitude.

Most manmade satellites that continue to function will exhibit some type of nonsmooth behavior or at least behavior that is less smooth than if only external forces were at work. They

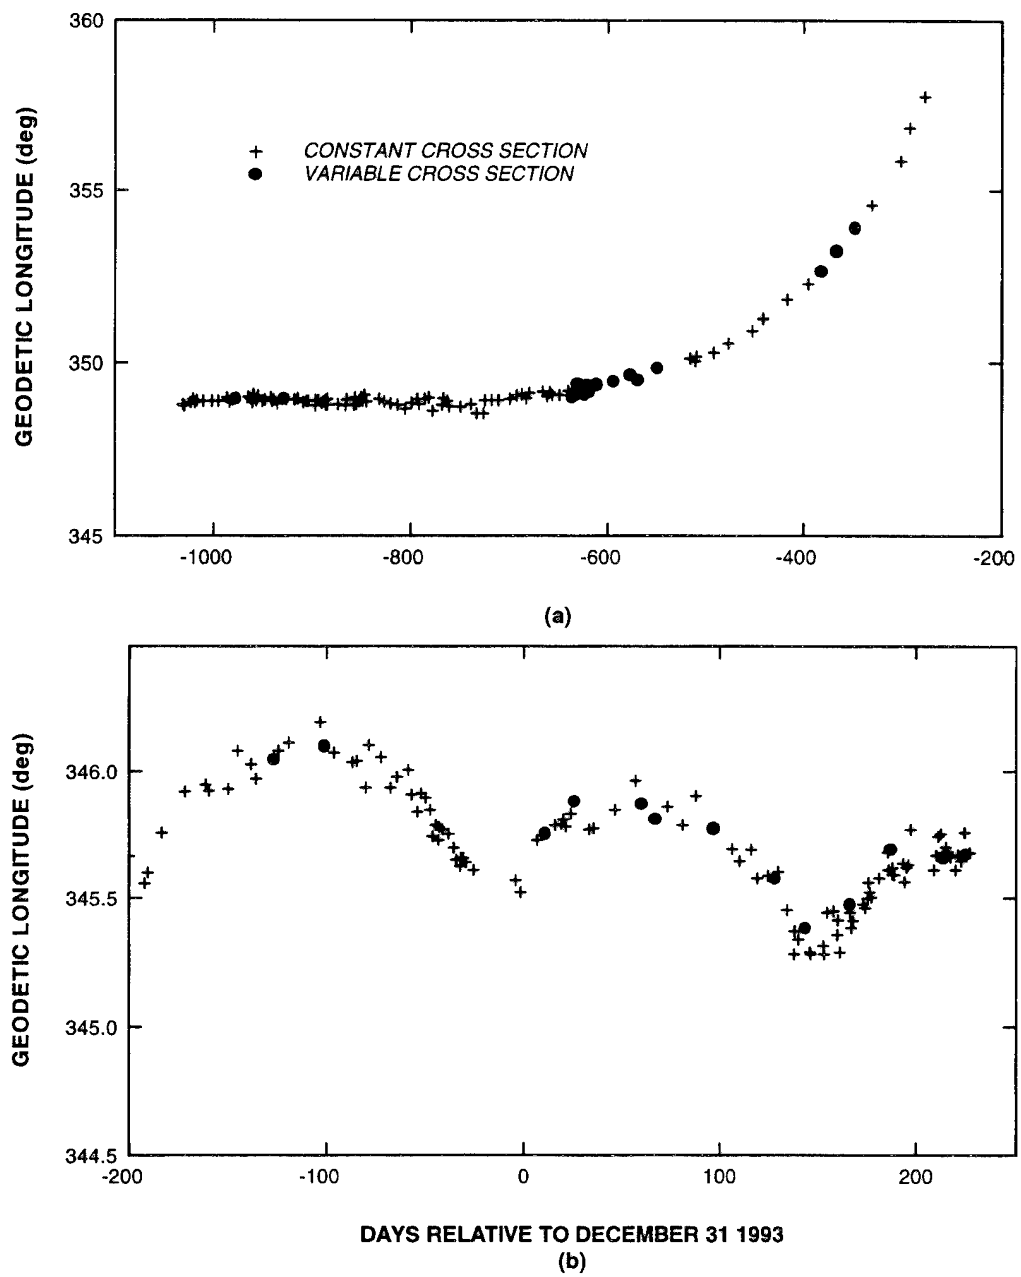

Figure 2. The orbital parameters of a satellite change smoothly from natural forces. These plots show the geodetic longitude of two synchronous satellites over time. In (a) the longitude smoothly increases with an increasing rate due to natural forces. In (b) the longitude changes nonsmoothly as thrust is periodically applied to keep the satellite near a particular value of geodetic longitude.

store energy internally, either chemically, mechanically or electrically, and release this energy to affect orientation or orbital motion. This may cause a net change in kinetic energy or convert kinetic energy between translation and rotation.

The nonsmoothness observed in the parameterizations of the orbits of satellites are fairly well understood. The main point of the remainder of this report concerns the role of randomness and nonrandomness in satellite status monitoring. This topic seems less understood by the space-surveillance community.

4.2 Random and Nonrandom Behavior

Natural changes in the kinetic energy of the satellite and the conversion of kinetic energy from one form to another follow complex laws. Although the orbital and rotational motion of the satellite can be modeled at some level, the use of the model requires precise knowledge of the structure of the satellite and the forces on the satellite. If we regard such knowledge to be unavailable, then the different dynamics evident in various satellites appear to be random. The rotation rates of some satellites increase while others have decreasing or constant rotation rates. Orbital energy decreases at different rates for different satellites, and can actually naturally increase. With enough knowledge most people believe these observations can be reconciled and predicted, but in our perpetual state of ignorance at least the fine details of satellite dynamics due to natural forces appears to be random.

Here is a specific example of how such randomness occurs. Suppose that a satellite maintains the orientation of its solar-panel array so that it is maximally illuminated by the sun. This means that the plane of the solar-panel array lies normal to the vector from the satellite to the sun. For such a satellite the value of the phase angle defined on page 11 roughly determines the set of possible viewing aspects from which the sensor views the solar-panel array.

Even with this constraint the cross section measured by the sensor is not always the same for a particular value of phase angle for several reasons. First, there is one additional degree of freedom in the orientation of the solar-panel array when its plane is aligned normal to the vector from the satellite to the sun; the solar-panel array may be rotated about the vector from the satellite to the sun. This still maintains the solar alignment of the panel but significantly changes the viewing aspect from the sensor. Second, other components of the satellite contribute to the cross section and these components, such as the main body and antennas, may be oriented separately. That is, the configuration of the satellite is unknown. Third, precise orientation of the panel is not critical so the panel might lie 10 or 20 deg from the ideal alignment. Any of these factors can cause significant change in the measured cross section, even though the phase angle is fixed and the solar panel array is oriented for solar illumination.

Thus the cross section appears to be a random process at a fixed phase angle for such a satellite. That is, the phase angle might be a good predictor of the average cross section, and the standard deviation of the cross section, but fails to predict the exact value of the cross section. Mathematically, it is said that for a fixed value of the phase angle the cross section is *random*.

Modeling the cross section as a random process over time at a fixed value of an independent variable, such as the phase angle, is the first key idea.

Most satellites that are still functioning also exhibit some form of nonrandom behavior with respect to some independent variable. In the previous example, the cross section is an example of random behavior at a fixed value of the phase angle, but the probabilistic description of the cross-section behavior (its mean and standard deviation) depends on the value of the phase angle. For one value of phase angle the measured cross sections have one particular mean and standard deviation. For another value of phase angle the measured cross sections have a different mean and standard deviation. Mathematically it is said that the cross section is a nonrandom function of the phase angle. Observing that the cross section is a nonrandom process over certain independent variables, such as the phase angle, is the second key idea.

There is an overlap of what might be regarded as nonsmooth with what might be called nonrandom. For example, the periodic correction of a satellite orbit to maintain the desired ground trace as a function of time is associated with both a nonsmooth and a nonrandom change in orbital energy. This can be seen in Figure 2, where not only do the orbital corrections cause a nonsmoothness in the geodetic longitude, but the corrections are applied periodically, in a nonrandom fashion. The remainder of the report is mostly concerned with characterizing the randomness observed in the behavior of satellite cross section.

For practical matters that are mentioned in Section 4.3, the amount of cross-section fluctuation is a more robust measurement than absolute cross section or the average cross section. The same arguments regarding the random and nonrandom behavior of cross section also apply to measures of the cross-section fluctuation.

4.3 Fluctuation in Satellite Signatures

For a complex satellite, the observed radar cross section is a complicated function of viewing aspect. As the viewing aspect slowly changes, a complex variation of the cross section is expected to be observed. For a given change in aspect angle the corresponding amount of cross-section change might be large or small. The cross section of a large solar-panel array is very sensitive to the viewing aspect when viewed nearly perpendicular to its flat surface, or one of its edges, and somewhat less sensitive at other viewing angles. The cross section of a large cylinder is very sensitive to the viewing aspect when viewed nearly perpendicular to its axis or its flat end surfaces and somewhat less sensitive at other viewing angles.

The cross section is obtained by a suitable scaling of the satellite's signature, so the amount of fluctuation in the cross section can be measured by measuring the amount of fluctuation in the satellite's signature. Section 5 describes how this can be done with the theory of runs. The measure of fluctuation, called a *runs statistic*, is a real number that can be computed for each satellite signature.

The discussion in Section 4.2 would suggest that the absolute cross-section values might serve as the basic measurement. Why not measure the absolute or average cross-section values instead of the amount of fluctuation in the cross section? The measurement of the fluctuation in cross section is less prone to error from radar calibration than the measurement of absolute cross section. This was discussed in Section 3.2. Measuring the amount of fluctuation in the observed cross section only requires that the radar remain stable for a few minutes. Its exact calibration state is not very important to the quality of the fluctuation measurement. Thus the fluctuation is more likely to be due to the satellite behavior than to the radar behavior.

If the satellite does not control its orientation, each time that its signature is sampled by the radar, a new and essentially unpredictable signature will result. This is not to say that the motion of the satellite does not follow precise laws of physics but that the external (and possibly internal) forces at work, the complex mechanical structure of the satellite, and its dynamic state at some reference time are all unknown. This means that the runs statistic exhibits apparent random values that do not depend on the values of the independent variables, such as the phase angle. There are interesting exceptions to this generality, and they are clarified in Section 4.4.

In Figure 3(a) the amount of fluctuation, as measured by the runs statistic, is shown as a function of the phase angle for a satellite that does not maintain its orientation or alter its configuration. The mean and standard deviation of the runs statistic appears to be similar at all values of the phase angle. The amount of fluctuation appears to be random at all values of the phase angle. Of course there are rigorous tests to confirm the hypothesis of randomness that are discussed in Raup[5], pages 23–30, and its references. In Figure 3(b) the fluctuation is restricted to a small range of phase values. This is a particularly obvious case where the amount of fluctuation depends nonrandomly on the value of the phase angle. It is believed that the satellite associated with Figure 3(b), which is maintaining its orientation, exhibits variable cross section near 0 and 180 deg of phase due to small aspect angle changes while the plane of its flat solar-panel arrays is nearly normal to the radar's line of sight. The increased fluctuation may be due to the sensitivity of the cross section to small changes in aspect angle when the plane of the satellite's solar-panel array is nearly normal to the boresight of the radar. Very small changes in aspect could cause large changes in cross section in this case.

If the orientation of the satellite is controlled, the randomness of the runs statistic is considerably reduced to a level consistent with residuals in the attitude control system and operational peculiarities. A satellite released from attitude control may not exhibit random behavior immediately. It takes time for its final controlled state to evolve into apparently random behavior.

4.4 Sources of Nonrandomness

The previous section might leave the reader with the feeling that any observed nonrandomness in the satellite characterization is due to the operational and structural peculiarities of an operating satellite, one performing some planned mission. This is not always true. Apparent nonrandomness

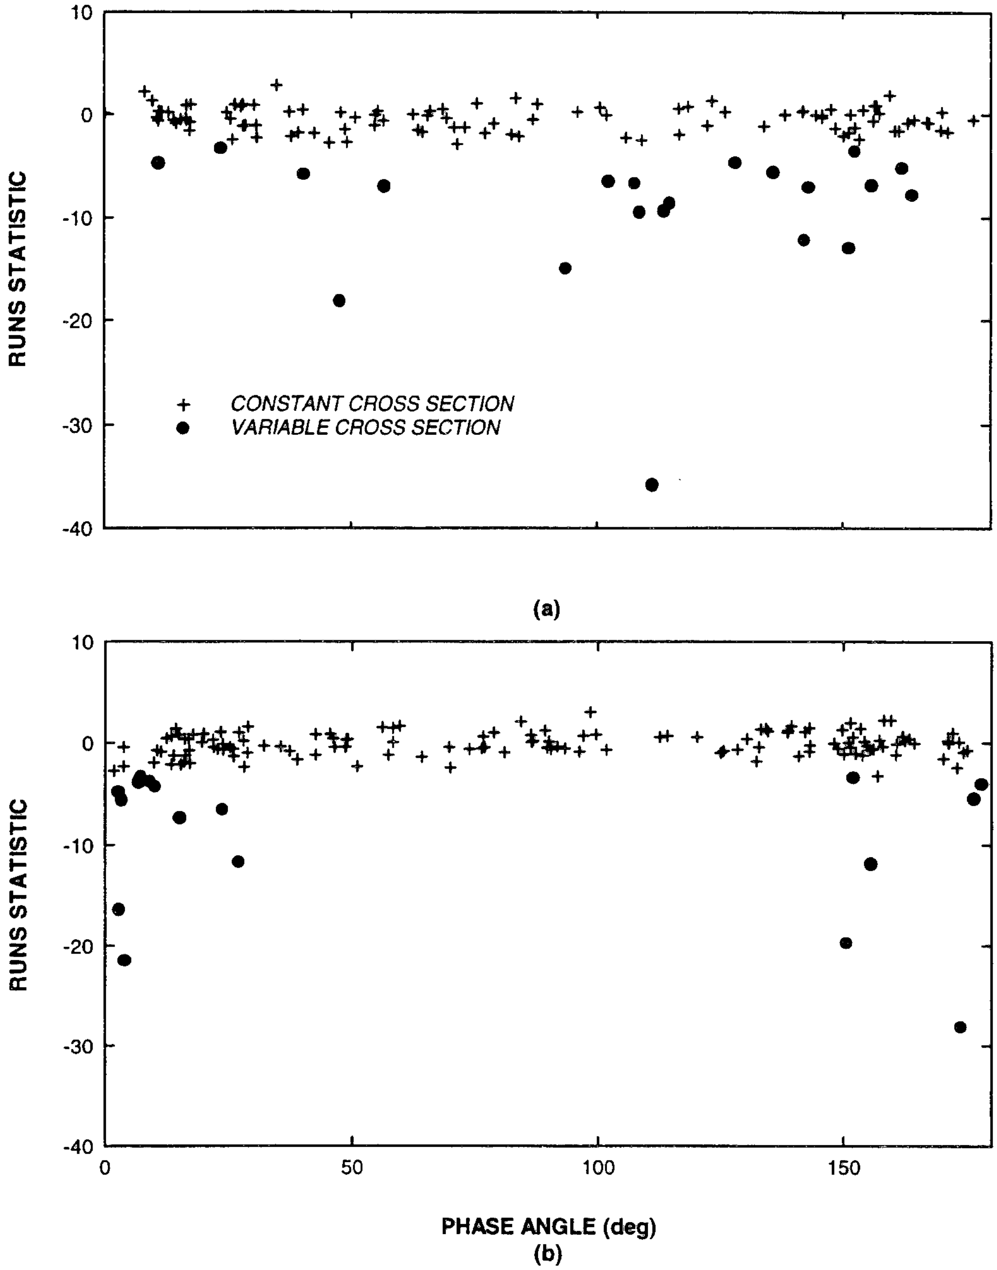

Figure 3. The amount of fluctuation observed in a satellite's cross section may appear to be a random function of an observation variable or not, depending on the satellite and which independent observation variable is chosen. These plots show the amount of fluctuation in the cross section of two satellites as a function of the phase angle. In (a) the observed amount of fluctuation appears to be random over all phase angles. In (b) the fluctuation is restricted to values of phase angle near 0 and 180 deg.

can result from the radar observation process. But once the source of observational nonrandomness is understood it becomes less of a problem for satellite status monitoring.

In Figure 4 a scatter plot of the binary classification of signatures as constant or variable cross section is presented as a function of the true anomaly and the standard deviation of change that

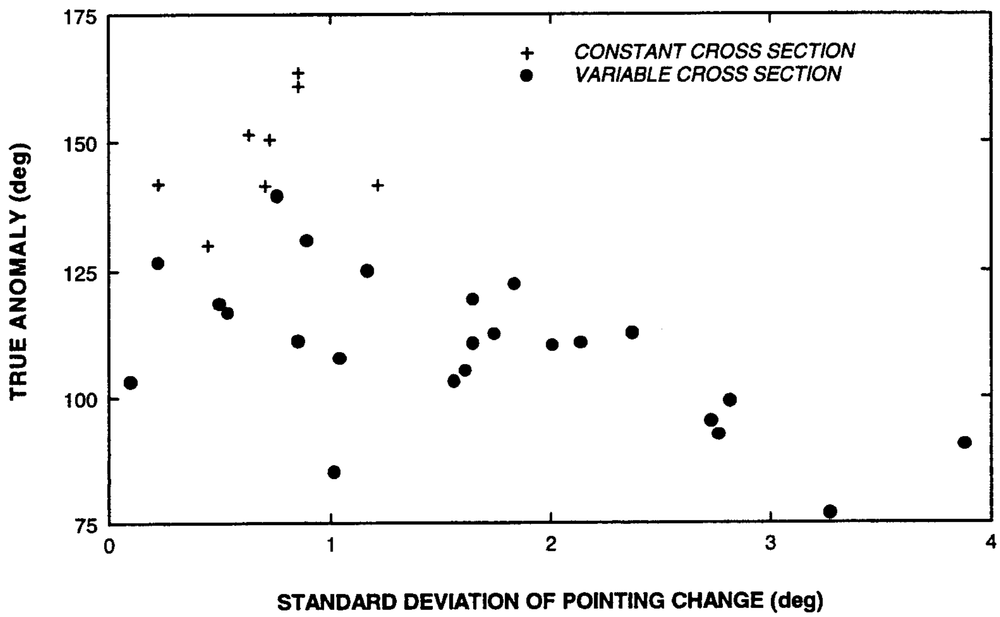

Figure 4. Nonrandomness in the observations of a satellite is not always due to the satellite's operational peculiarities. The clustering of constant cross-section signatures in the upper left-hand corner arises from the observation process and not from nonrandom behavior of the satellite. Once the source of observational nonrandomness is understood it becomes less of a problem for satellite status monitoring.

occurred in the radar pointing while observing the satellite. The change in radar pointing occurs because the satellite is moving along its orbit, relative to the radar. If the satellite were inertially stable, the amount of change in the radar pointing would correspond to the change in the aspect angle from which the radar observes the satellite. It is helpful to think of the satellite as fixed in this case, with the sensor moving through an angle corresponding to the change in pointing. Thus some amount of cross-section fluctuation may be due to the change in the sensor pointing, or equivalently, the motion of the satellite along its orbit.

The cluster of constant signatures in the upper-left corner of the graph is due to the observation process, not nonrandomness of the satellite behavior. This satellite is nearly inertially stable. It rotates very slowly under the influence of a natural torque. When the pointing angle change is small, there is very little change in the observed cross section and the signature often appears to be constant. Furthermore, as explained in Section 5, the classification of the signature as constant or fluctuating is done by an algorithm with a constant probability of erroneously classifying a constant signature as fluctuating. When the signal-to-noise ratio is small, as when the satellite is near apogee at a true anomaly near 180 deg, then the algorithm has a small probability of detecting a fluctuating signature. The signal-to-noise ratio decreases as the satellite moves farther from the radar. As the signal-to-noise ratio decreases more fluctuation must occur before the algorithm can classify the signature as fluctuating. The relative stability of the satellite, small pointing angle changes, and the increasing range to the satellite all contribute to the clustering of constant signatures that suggests a nonrandom dependence of the amount of signature fluctuation on true anomaly.

It is probable that nonrandomness can arise from other observational and natural factors. As yet, other concrete cases have not been identified, but the reader should realize that the causes of observed nonrandomness in satellite characterizations may be complex and sometimes obscure.

4.5 Characterization Summary

Because the narrowband radars of the SSN cannot be used to directly extract the physical size, shape, and orientation of target satellites, some other approach is needed to apply the radar data to the satellite status monitoring problem. An indication of whether or not the satellite is performing some useful mission follows from noting the smoothness and randomness of its behavior. The remainder of the report exploits the idea of randomness in its behavior.

The normal operation of the satellite may constrain the orientation of solar panel arrays, antennas, sensor heads and the main body of the satellite. The constraints will be related to independent variables such as the phase angle, apparent elevation of the satellite as seen by the sensor, time of day, geodetic coordinates, and the true anomaly of the satellite. Thus the cross section of a satellite may appear to be a nonrandom function of one or more of these variables if the satellite is being used to perform some useful function. At a particular value of these variables, the imprecise nature of the constraints causes the cross section of the satellite to vary randomly.

A satellite might be characterized by collecting cross-section measurements restricted to certain narrow ranges of the important independent variables where the cross section behaves randomly, parameterizing the observed randomness of the cross section and monitoring for a change in these parameters. The system will be less sensitive to radar calibration problems if the amount of fluctuation in the satellite cross section during the observation is used as the basic observable instead of the average cross section, which is subject to bias errors.

These ideas will be made more precise in Section 5 as the algorithms for implementing the satellite characterization scheme are outlined. Also in that section, algorithms that use the characterization to flag unusual observations and potential changes in the satellite characterization are developed.

5. SIGNAL AND DATA PROCESSING ALGORITHMS

This section describes the signal and data processing algorithms used in the system. The algorithms provide satellite characterization and anomaly detection using the principles illustrated in Section 4. This section provides an intuitive description of the algorithms, the mathematical basis of which can generally be found in Raup [5]. The software architecture of the system's implementation is deferred until Section 6. An overview of the processing described in this section is provided in Figure 6.

5.1 Signal Processing

The facility must have one or more signal processing components. Signal processing normally resides on a computer at the sensor site, where it effectively compresses the observational data into a form suitable for a narrow-bandwidth transmission. In compressed form the observation requires much less than a kilobyte of storage. This compression is important when data processing is done at another location. The data processing component of the system needs only the compressed results of the observation, not the raw observation itself. The observation collected by the sensor consists of target signature data (that can be used to estimate the satellite's cross section as a function of time), an orbital element set for the satellite, and sensor state information (such as pointing angles from the sensor to the satellite and so forth), all as a function of time. These data are stored at the site for a period of time. The signatures are also smoothed and stored in compressed form for transmission on demand to the analysis interface component of the system where they can be displayed by the analyst. Typically much less than one percent of signatures are of interest to an analyst.

The major signal processing operation provides the basis for classifying signatures as either constant cross-section signatures or variable cross-section signatures as follows. When a satellite exhibits a constant cross section versus time, the signature observed by the sensor is ideally a stationary random process determined by the constant power level due to the target and the constant power level of random observation noise. When the cross section of a satellite fluctuates, its signature is no longer mathematically stationary because the power level due to the target cross section varies, causing changes in both the mean and higher order moments of the signature. A test for stationarity can therefore be used to test whether the signature is associated with constant or variable satellite cross section.

The signal processor uses a *runs test* as described in Raup [5], pages 5–9. Briefly, the test consists of computing a number called the runs statistic that is compared to a threshold determined by implementing a standard two-tailed binary hypothesis test. Small and equal probabilities are applied to the tails of the density for the runs statistic under an assumption of stationarity to establish upper and lower thresholds for testing the value of the statistic. If the value of the statistic falls between the upper and lower threshold then the hypothesis of stationarity is accepted and the signature is declared to be of the constant cross-section type; otherwise, it is declared to be of the

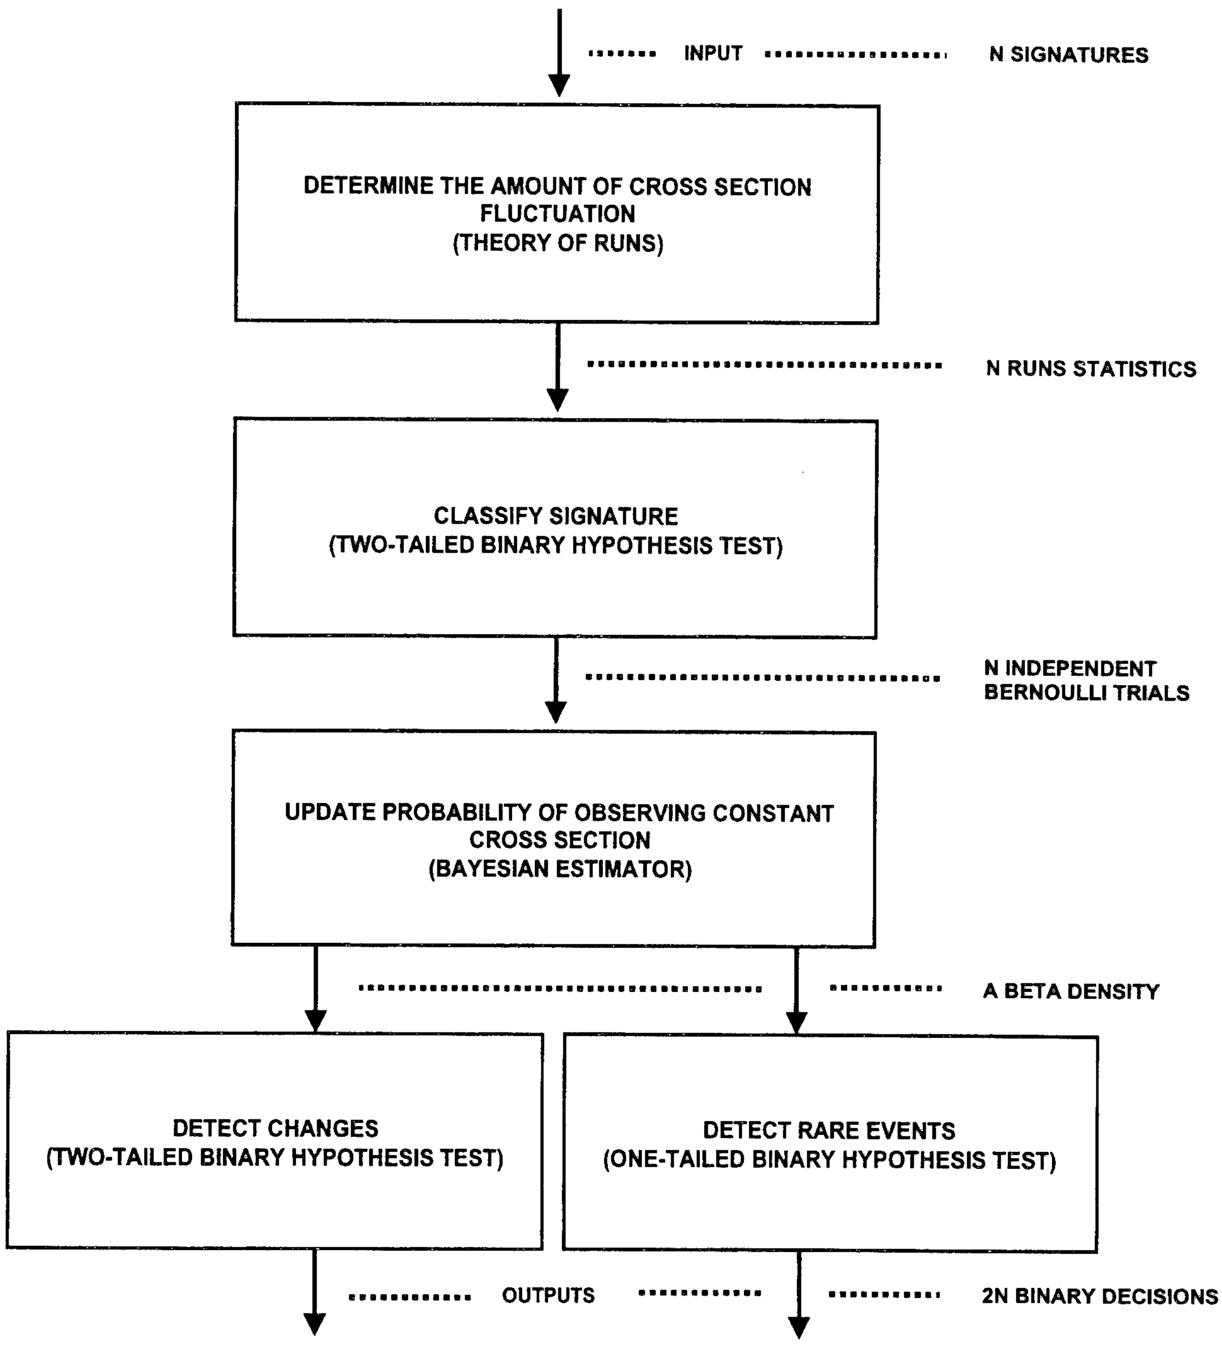

Figure 5. A variety of statistical and estimation-theoretic methods are employed in the signal and data processing components of the system. The diagram outlines the steps required to transform N signatures into $2N$ binary anomaly detection decisions.

variable cross-section type. Selecting the thresholds in this fashion implements a simple constant false-alarm test. The probability of erroneously declaring a constant cross-section signature to be variable is equal to the sum of the two-tail probabilities. This small value of probability is called the significance of the classification test, and one minus its value is called the confidence of the classification test.

The probability of detecting variable cross-section signatures must then vary in a way that is dependent on the data. If the amount of fluctuation in cross section is very small, or if the signal-to-noise ratio is very small, then cross section variability may go undetected. This is a property of all constant false-alarm rate algorithms. It may mean, for example, that a satellite may consistently exhibit variable cross-section signatures that are obvious at perigee but the variability of the cross section may go undetected at apogee where there is much more noise in the signature.

The runs statistic, before it is compared to the threshold for the classification test, provides a measure of the amount of cross-section fluctuation in the signature. The runs statistic can be used to objectively order signatures by the amount of cross-section fluctuation, subject to the signal-to-noise limitations discussed above. The runs statistic stored in the signal processing data base is a *normalized* statistic, which helps to remove some of the effects of the observation duration, or equivalently the number of signature samples used to compute the statistic.

The value of the statistic will be near zero for a constant cross-section signature. Its magnitude will increase as the cross-section fluctuation increases. Generally the runs statistic will become more negative. Its values range from about $+3.0$ to -60.0. Although the thresholds for declaring variable cross-section signatures depend on the value chosen for the probability of classification error, the analyst will observe that runs statistic values below about -3.3 or above about 3.3 will cause the signature to be classified as a variable cross-section type.

One last aspect of the problem of measuring the amount of cross-section fluctuation with a runs statistic needs to be considered. As mentioned, the statistic will have a value near zero for a constant cross-section signature. As the amount of fluctuation increases, the value of the statistic will become more negative. If the amount of fluctuation continues to increase, the bandwidth of the signature will exceed the sampling bandwidth of the sensor, a conditioned called *aliasing* in sampling theory. As aliasing increases, the signature becomes mathematically more stationary once again and the value of the runs statistic will become more positive, eventually becoming near enough to zero to cause the system to believe that the signature is a constant cross-section type. For this reason it is possible for the cross-section fluctuation of a rapidly rotating satellite to go undetected when its signature is extremely aliased.

The signal processing system also applies the runs test to the time series that characterizes observation noise in the sensor. These noise sequences are selected so that they become nonstationary when certain tracking problems occur that could adversely affect the quality of the signature data. Tracking problems can cause the observed cross section to fluctuate even though the satellite cross section is constant, and it is undesirable to erroneously assume that such fluctuation is due to an attribute of the satellite. The data processing component of the system will not use data that are associated with nonstationary noise sequences. The techniques for rejecting bad satellite observations are discussed in more detail in Section 8.

5.2 Data Processing

Although the signature data is stored with the signal processing component of the system, only a few hundred bytes of signal processing results need to be routinely sent to the data processing component of the system. This works well when the results from several signal processing components (and presumably several sensors) must be fused by a data processing component at another location. The signature data itself is only sent when requested by an analyst; in practice less than one percent of such data is really needed in raw form.

These few bytes summarizing the signal processing results, along with the orbital element sets and the sensor state information, less than one kilobyte, are sent to the data processing component of the system. Data processing is used to extend the basic observational data, fuse the signature classification results with the historical characterization of the satellite, and flag anomalous observations. These operations are now described.

First the orbital element set is used to predict some of the sensor state information that came with the observation. In fact, the predictions are compared with the information reported by the sensor to determine if the observational data is consistent with the orbital elements. Currently the azimuth and elevation pointing state reported by the sensor is compared with orbital predictions. Tolerances for allowable discrepancies are set in a system file. The sensor-supplied runs test results for the noise sequences are also evaluated. If any of the consistency tests or stationarity tests fail, then the data is not reliable and is not used.

5.2.1 Satellite Characterization

One of the major functions of the data processing component of the system is to fuse the signature classification results into a characterization of the satellite. It does this by maintaining an estimate of the probability of observing a constant cross-section signature for the satellite when viewed under certain observation conditions.

What is meant by the probability of observing a constant cross-section signature? Signatures are classified as being either a constant cross-section type or a variable cross-section type, like a coin tossed onto a table reveals either a head or a tail. Unlike the coin example, where the probabilities of seeing the head or tail are both 0.5, the probability of observing a constant cross-section signature is generally not 0.5 and is unknown. Its value varies from satellite to satellite. For a single satellite its value changes over time and for the different geometric conditions under which the satellite is observed.

The data processing component of the system uses the signature classification history of the satellite under a restricted set of observation conditions to estimate the unknown value of the probability that the next observation of the satellite will reveal a constant cross-section signature.

Uncertainty about the true value of the probability is represented by a beta probability density function. Specifically, the data processing system feeds the signature classification history of the satellite into a Bayesian estimator for a beta density. This density is the mathematical object that characterizes the satellite based on past observations. The mean of the density (called a conditional mean) is used to predict the probability of observing a constant cross-section signature, and a confidence interval for the prediction is obtained by eliminating small and equal probabilities from the tails of the density. Details of the algorithms are found in pages 11–16 of Raup [5].

The observation conditions that are outlined above are currently part of the system architecture. They do not change except through parameter settings that can be altered by the administrator. They are required because the characterization of the satellite can vary as a function of viewing geometry, orbital position, time of day, and so forth. Satellites are currently characterized throughout a range of phase angles (the angle between the vectors from the satellite to the sensor and from the satellite to the sun) that exclude the case that the sensor, the satellite, and the sun lie nearly along a straight line, a condition called *solar alignment*. Satellites observed near solar alignment often exhibit variable cross-section signatures presumably from small-angular motions coupled with the large cross-range extent presented to the sensor by solar panel arrays. Satellites in highly eccentric orbits are characterized near their apogees, because they often exhibit variable cross-section signatures due to high angular tracking rates at other positions in their orbit. By excluding observation conditions where nearly all satellites in all phases of operation exhibit variable cross-section signatures, the characterization of the satellite becomes more sensitive to changes in the satellite status and operation. Thus when examining the satellite characterization results, some otherwise good sensor observations will have been excluded. Observation conditions in the context of a particular database from a particular sensor are discussed in more detail in Section 8.

5.2.2 Anomalous Observation Detection

The second major function of the data processing component of the system is to flag anomalous observations of the satellite. It does this by applying two algorithms designed to find inconsistencies between the new observation and the beta probability density function that serves as a characterization of the satellite. One algorithm uses the density to flag signature types, either constant or variable, that occur less than a specified fraction of the time. This algorithm is a rare signature classification event detection algorithm. The other algorithm examines recent signature classification results to identify changes in the underlying probability of observing a constant cross-section signature. This second algorithm is a satellite characterization change detection algorithm. The operation of both algorithms are explained in the following paragraphs.

5.2.3 Rare Signature Classification

For some satellites a variable cross-section signature does not occur very often. For other satellites the same is true of a constant cross-section signature. For still other satellites neither variable nor constant cross-section signatures are unusual.

The data processing system flags rare signatures, those types (either constant or variable) that occur less than a certain fraction of the time. Details of the test are found in Raup [5], pages 17–21. For argument's sake say the value of the fraction is 0.2, a typical value. If a classification type occurs less than this fraction of the time it is called *rare*; the other classification type is called *common*. A one-tailed test is applied to the beta density for the probability of observing constant cross section in order to determine if either type of signature classification is rare. In this example the hypothesis that constant cross-section signatures are rare is accepted if all of the probability mass of the density, except for a very small part, lies between 0 and 0.2 probability of observing a constant cross section. The hypothesis that variable cross-section signatures are rare is analogously tested.

The value of the small probability mass mentioned in the previous paragraph is the probability of erroneously declaring constant cross-section signatures rare when they really are not, hence it is set very small. Fixing this probability for the tests means that the probability of detecting rare signature classifications varies, depending on the available data. For reasonable probabilities of error the data processing system must collect about 13 signatures from a very stable satellite before it recognizes that variable cross signatures are rare, for example. The value of the small probability mass is called the significance of the rare event test, and one minus its value is called the confidence of the rare event test.

The existence of rare signature classification events does not necessarily mean that anything has happened to the satellite, but they are often clues to the behavior of a satellite. For example, the rare events may tend to occur at certain times of day for synchronous satellites, possibly related to an aspect of the satellite operation. A run of several consecutive signatures that causes a change detection flag often begins with a rare event, so it is possible to get advanced warning of a change detection by examining certain rare events. The change detection algorithm is described next.

5.2.4 Satellite Characterization Change

While the beta density for the probability of observing constant satellite cross section is being estimated from the signature classification results, an assumption is made that the unknown probability is constant. After all, if the probability were to change while the density was being estimated, then the old estimate should be discarded and the process of building a new one should begin. The data processing component of the system applies a change detection algorithm to the new signature classification results as they arrive from the sensor.

The algorithm works by testing run lengths and causing the partially constructed probability density to be discarded if the observed run lengths are inconsistent with the density. In this context a *run* is one or more consecutive signature classifications of the same type. This concept of a run is related to the statistic used by the runs test for classifying satellite signatures, but it is defined differently for this change detection application. Details of the test may be published in a future technical report, but at this time a detailed description is not generally available. Suppose five constant cross-section signatures are followed by three variable cross section signatures. This event event is recorded as

$$\ldots \mathrm{CCCCCVVV}$$

and note that the observation ends with a run of three variable cross-section signatures.

The data processing system transforms the beta probability density for observing a constant cross-section signature into the probability densities for the lengths of runs of constant and variable cross-section signatures. A two-tailed test is constructed by choosing small and equal probability masses in the two tails of the appropriate density to establish two run-length thresholds. If the observed run length is shorter than the smaller threshold, or longer than the larger threshold then the beta density characterizing the satellite is rejected, and its construction begins again.

An analogous example occurs when tossing coins. If in tossing a new coin 50 times one were to observe 50 heads, then the experimenter might be led to believe that closer examination of the coin would reveal that it had two heads! He or she would reject the notion that the probability of flipping a head was one half. However, if there was a good mix of heads and tails, and only the last 3 of the 50 flips were consecutive heads, then all would seem right with the universe.

The sum of the tail probability masses that determine the run-length thresholds is the probability of erroneously rejecting the beta density when it is actually the correct one, so the probability masses must be small. Fixing this probability means that the probability of detecting changes in the underlying probability of observing a constant cross-section signature varies, depending on the available data. The value of the sum of the tail probability masses is called the significance of the change detection test, and one minus its value is called the confidence of the change detection test. If the change in probability is small, or there are not many observations of the satellite then the change could be missed. This type of performance characterizes all constant error-rate algorithms.

A change detection does not mean that the satellite has changed, only that the recent data are inconsistent with what is known of the satellite characterization so far. Change detections can be triggered by changes in when, where, and how the sensor collects data. They can also be caused by bad sensor data. However only about 1% of signatures trigger change detections overall. This percentage is smaller for some satellites and larger for others. A fraction of the change detections are associated with real changes in the satellite behavior, so investigating each one often leads to a discovery about a change in the stability or behavior of a satellite. An example is provided in Section 7.

6. SYSTEM DESIGN CONSIDERATIONS

The system consists of the components illustrated in Figure 6. All of the components except the sensor-dependent data server are included in a portable software distribution package for UNIX platforms that consists of *tar* files, installation instructions, and an interface document. Every sensor produces data in its unique format. Therefore, the interface document is provided as a guide for converting a sensor's particular data formats to the standard input forms for the system. Without the interfacing convention the system software would not be portable.

The system uses distributed computing across multiple UNIX platforms to place its signal processing, data processing, and analysis components at the appropriate geographical locations. Of course any two or all three components may reside on a single UNIX machine, if desired. The system must have one or more signal processing components that archive the sensor data and compress it for transmission to a centralized data processing facility. The signal processing component should be placed near the sensor site because of the fairly large bandwidth required to transfer data from the sensor. The data processing component handles all satellite characterization and anomaly detection. Interactive access to the observation database is accomplished through one or more analysis interface components that may be located far from the centralized data processor. The analysis component includes a graphical user interface (GUI) based on X Windows with Motif. Data networking is accomplished via network file system (NFS) mounts. The GUI networking is built into X, including support for PC compatibles and Macintoshes.

Throughout the design UNIX built-ins were used as much as possible to minimize the amount of code written for the project. When a UNIX built-in was not available, a processor was written in Fortran or C as required. Most scientific processing has been done in Fortran, partly to directly access Fortran scientific processing libraries. The GUI residing in the analysis interface component uses much C code in order to directly access the X Windows and Motif libraries. The UNIX built-ins and custom executables compiled from Fortran and C source code are integrated with Bourne shell scripts.

Each component of the system maintains its own local databases. The databases are ASCII text so that they are portable and so that the many text-based UNIX built-ins can be used to manipulate them. The scientific database files are generated and updated with the *make* utility. The processing required to generate and update the scientific databases are documented as rules in the *makefile*. Although *make* was probably not originally designed as a database maintenance tool, database generation and maintenance are natural applications for *make*. The interdependence of the various elements of the databases and incorporation of new measurements into the proper database elements are naturally handled with the *make* paradigm.

Extensive use is made of stream editing techniques for data manipulation. The stream editor commands appear as rules in the *makefile* for database generation and maintenance. Most of the stream processing is implemented with *sed*, *awk*, or custom code written in Fortran or C. The signal and data processing components, where most of the scientific processing is done, uses most

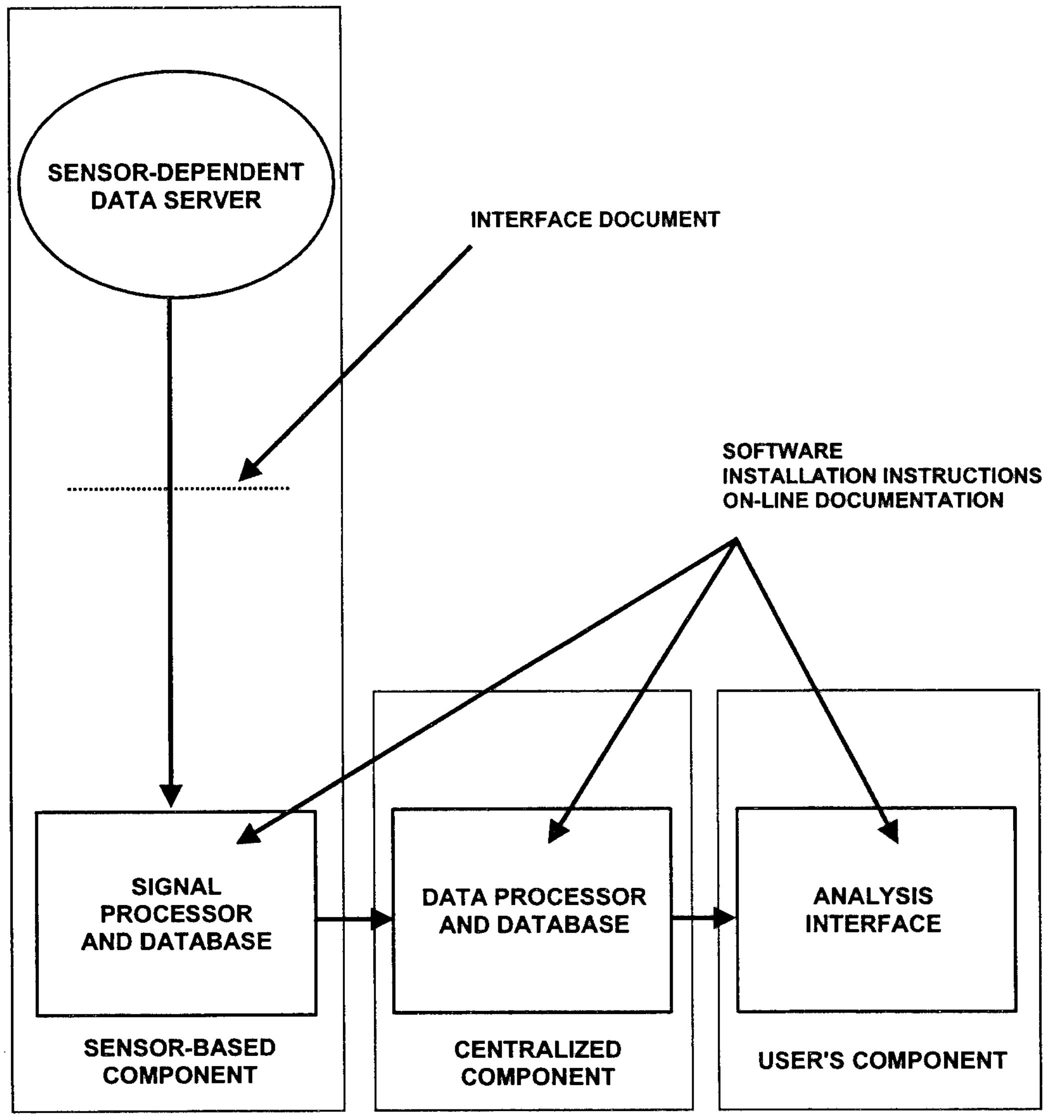

Figure 6. The system is a portable network-transparent UNIX software design that is connected to one or more sensors through a documented standard interface.

of the Fortran code. The analysis interface uses most of the C code because of the GUI based on X Windows and Motif.

Only generously licensed software is used in the system in order to avoid the cost and maintenance issues involved with proprietary software. X Windows and Motif make an excellent portable base for the GUI, and government and university code provide excellent libraries for scientific data processing.

Section 6.1 describes the data server. This is software conforming to the interface document that acts as a translator between the sensors, which are all different, and the system, which accepts data in the form of standard files. The signal and data processor implementations are discussed in Section 6.2. The analysis interface component is discussed separately in Section 6.3. The system was designed so that certain classes of new analysis tools can be easily added. This is discussed in Section 6.4.

6.1 Data Server

The data server has a sensor-dependent design. Its purpose is to convert the data from the unique form peculiar to the sensor to the standard form used by the portable signal processing component of the system. The interface document that is included with the software distribution specifies the output of the data server. Because of the uniqueness of each sensor, the data server is not portable and it is not included with the software distribution. This section describes some features of the Millstone Hill data server that connects the L-band and UHF radars to the developer's system. Even though portable software could not be written, some of the features discussed in this section may be useful when designing data servers for other sensors.

The data server has been implemented at Millstone Hill as a stand-alone UNIX program. The data server processes raw data files produced by the program responsible for real-time process control of the radars. These two programs reside on separate machines attached to the same local ethernet in a secured environment. The radar control computer delivers the files via an NFS mount onto the appropriate file system of the data-serving computer.

The main function of the data server is to produce the standard files described in the interface document that is included with the software distribution. The server also is tasked to distribute standard files to any number of other machines with the signal processing component of the system. When the other machines are attached to the same ethernet as the data server, then NSF mounts are used to deliver the files. When the machines are not on the ethernet attached to the data server, then *tar* files on magnetic tape are used.

The data server executes periodically, processing any new data provided to the data server via the NFS mount described above. One of the goals in implementing the data server on the classified network at Millstone was to provide the test director with near real-time feedback from the system. The test director is responsible for operating the radar system. The design of the data server currently allows for feedback to the test director within a few minutes of a successful track.

Executing the data server causes the following tasks to be performed for every sensor data file provided to the server.

1. The sensor data file is opened and read to determine the *observation key* as described in the interface document.

2. The sensor data file is renamed based on the observation key.

3. The sensor data file is processed through the program that generates the standard files.

4. The standard files are copied to the input directory of any signal processing components installed on the same machine and to any export directories for signal processing components installed on other computers attached to the same local ethernet.

5. The standard files are moved to an archive destination based on security classification of the data.

The archive destinations mentioned in the last step above are two distinct disk directories. Classified files are written to one directory and unclassified files are written to the other. Archiving of the data is performed in two steps. The unclassified data are archived once a day by copying the contents of the unclassified archive directory to magnetic tape using the *tar* utility. The classified data are archived once a month by copying the contents of the classified archive directory to magnetic tape using the *tar* utility.

The archival process provides a method for moving standard files to machines attached to different ethernets. At Millstone Hill there are two networks: one classified and one unclassified. Distribution of standard files to classified computing systems can easily be handled via NFS mounts. Distribution of standard files to unclassified-computing systems is achieved utilizing the daily unclassified archival tape. Standard files are copied from the unclassified tape to the input directory of the signal-processing component of the system on the destination machine.

The construction of the unclassified tape occurs on a track-by-track basis. For every sensor data file there will exist a set of standard files associated with some observation key. These standard files are always written to tape so that the token file is unarchived last, as prescribed by the protocol of the interface document.

The system applies simple tests to the data found in the standard files in order to reject bad observations, but some data editing within the data server may be appropriate. Each sensor has its own peculiarities that are well-known by site personnel. For example, at Millstone Hill it was not uncommon for bad pointing data to find its way into the data stream, which made the calculations of the standard deviations of pointing invalid in downstream processing. A simple algorithm was implemented to edit the bad pointing data from the standard file.

Also Millstone Hill's processing does not consistently supply both principally and orthogonally polarized receiver data to the server. Sometimes the orthogonally polarized data are missing. If both polarizations are always available, then the best course of action is to combine them into a single signal-plus-noise channel. When both are not available only the principally polarized channel data can be used to estimate power. Because switching between these two methods might make the downstream processing results incomparable, it was decided that the orthogonally polarized data would not be used within the Millstone Hill data server.

6.2 Signal Processing and Data Processing Components

The introductory remarks at the beginning of this section provide an overview of the design considerations applied to the signal-processing and data-processing components of the system. These components perform most of the scientific processing associated with the system using a variety of statistical methods. For further details the reader is referred to Section 5.

6.3 Analysis Interface Component

The introductory remarks at the beginning of this section provide an overview of the design considerations applied to the analysis interface component of the system. In this section the system architecture of the graphical user interface included with the analysis interface is considered in more detail. The GUI is based on X Windows with Motif. The analysis interface is designed to provide a flexible data analysis capability. In order to meet this goal the system architecture includes two unusual features. The first is to use scientific data plots not only as a means to visualize the data but also as a means to perform simple relational data selection without a formal relational database engine. The second is an interclient communications package that allows user interaction with one analysis tool to alter the performance of other analysis tools executing as separate UNIX processes.

From the user's point of view the GUI is documented by its built-in help functions as well as more detailed UNIX manual pages. All of the user commands in the analysis interface are implemented as stand-alone UNIX executables, called analysis clients in this report. Most of the analysis clients, but not all, are also X clients. That is, they display one or more windows in the X environment. Because the interface is composed of standard UNIX and X Windows programs, other standard programs such as word processors, calendars, clocks, and email tools may be used with the analysis interface. Typically, the interface includes an instance of *xterm* so that the analyst may enter standard UNIX commands. Although a knowledge of UNIX is not required to operate the interface, all of the power of UNIX is available concurrently with analysis interface operation.

The GUI included with the analysis interface is implemented using the X Windows graphics system for several reasons. X is the most common window system and is available for a wide variety of machines, including UNIX workstations, IBM-compatible PCs, and Macintoshes. X also provides built in networking capability. Thus, the user interaction can take place from any computer with an X server while the actual analysis interface programs reside on a remote UNIX machine. X provides a flexible, network-transparent base upon which the interface is built, allowing it to use a variety of client-server configurations with no extra effort from the user or programmer.

The analysis interface includes three main X clients and an interclient communications package. Various instances of the X clients are used to display information and to interact with the user. The interclient communications package allows the analysis clients to share information as the state of the analysis interface changes.

6.3.1 X Clients

The three main X clients are *xcomut* and *dyninfo* (developed specifically for the system) and a modification of the generously licensed *xgraph* program. Throughout this report any mention of *xgraph* generally refers to the version as modified for the system. The program *xcomut* allows the user to launch and manage other analysis clients. The program *dyninfo* provides textual information that is not easily or naturally incorporated into a graph. The program *xgraph* is a scientific graphing program that allows the user to see and to select relevant data with the mouse. Each of these programs is a standard UNIX and X Windows program. When they are used together, they interact to provide the data exploration functionality of the analysis interface. In any one analysis session the user may have many instances of *dyninfo* and *xgraph* active in order to display a variety of text and scientific data plots. Typically, only one instance of *xcomut* is present in an analysis session because it is capable of managing an arbitrary number of analysis clients.

The Motif tool kit has been used to implement *xcomut* and *dyninfo*. Motif is a library of routines for incorporating sophisticated graphical user interaction into a program. It provides the "look and feel" of the graphical interface and easily allows for extensive customization to suit user preferences. Using Motif does not require the analyst's computer to have any additional software. Most importantly, Motif will allow the functionality of the interface to be easily modified or extended in future versions of the software.

Xcomut (CommandUtility) An instance of the program *xcomut* serves as the *CommandUtility* for the analysis interface. The *xcomut* program is a flexible Motif client that can be used to start any type of analysis software that is written as a stand-alone UNIX executable. It provides consistency by allowing the user to start data analysis clients with simple, intuitive mouse clicks, replacing complicated textual commands. It provides a central display of information outlining the functionality of each client. The graphical nature of *xcomut* and its simple interface make it an efficient way for the analyst to control the analysis environment.

When a *CommandUtility* is started, a user is presented with a list of available analysis clients. After selecting a client from the list with the mouse, the user is given information about that client. This information includes a description of the analysis client, a list of command-line arguments that can be edited, the ability to get further help on the command-line arguments, and a string that can be edited and is suitable for launching the client from the command line. The *xcomut* program uses the values of the interface environment variables as defined in Section 6.3.2 and the client's instance definition as defined in Section 6.3.5 to determine the default arguments for launching clients. The user may edit the default argument values or add new arguments directly to the command-line display. The user may then use pull-down menus to launch the client in either static or dynamic mode. Dynamic mode is only available when the client is specifically written to include support for it. The pull-down menus also allow the user to get help in using *xcomut*, to kill all dynamically launched clients, and to quit *xcomut*.

To review and execute a command with its default arguments, the user only needs three mouse clicks. An experienced user has many other options. The values in the command-line argument boxes of the *CommandUtility* display may be changed by the user. When the user changes a value, the change is reflected in the command-line string. The change does not affect the interface environment variables or other clients. The user may also make changes directly to the command-line string, but these changes will be lost if the user changes the command-line arguments after changing the command-line string. When the user launches a client, either statically or dynamically, the value of the command-line string, including any user modifications, is used to launch the client. Because the command-line arguments typically depend upon the values of the interface environment variables, the ability to change them within the *CommandUtility* before launching gives the user the option to launch clients with different, local values of the interface environment variables. Editing the command line directly allows the user to add arguments that normally aren't used by the interface, such as the standard X Windows arguments for controlling window appearance.

Dyninfo The *dyninfo* program is a Motif client for displaying text. It is used in various instances to summarize data about a satellite or a single observation of a satellite. It is used in the analysis interface to display information that does not easily or naturally fit into a coordinate graph. The program *dyninfo* may be launched dynamically so that the information displayed remains current when the user interacts with other clients or statically so that needed information remains regardless of the user's interaction with other clients.

Xgraph The *xgraph* program is scientific plotting software originally obtained from the University of California that can display arbitrary line and scatter plots with a variety of colors and markers. It is used in various instances to display sets of data pairs (scatter plots) and functions of a single variable. The analyst can use the mouse to obtain increasingly enlarged views of areas of the graph. Although there are a variety of scientific plotting programs available for X Windows, none had exactly the features that were needed. The *xgraph* program was chosen because it had most of the capability that was needed, and it was still simple enough to easily modify.

It is common to display functions of a single variable as a graph in a scientific plotting utility but it is less common to use the graph to select interesting data points. The principal modifications to *xgraph* were made to enable it as a means of selecting data with the mouse. In addition to one or more data files that result in displayed graphs, the modified *xgraph* program reads an action file. This file defines regions of the screen surrounding the displayed data point that become sensitive to the mouse. When the analyst clicks a mouse button in a sensitive region of the screen, information read from the action file is passed to the interclient communications package.

6.3.2 Interclient Communications Package

User interaction with the analysis interface changes its state. The display state is described principally by the values of the interface environment variables. The variable values are stored in a file as strings of the form *variable* = *some value*, one string per record. The file containing the values of the interface environment variables is part of the interclient communications package. The interclient communications package provides the means to alter the state of the interface by changing the values of the interface environment variables. As the state of the interface changes, the interclient communications package is also used to determine how to launch the clients that implement the user commands and what processing each client should be doing. The interclient communications package consists of various executables, information files, and a protocol describing their proper use. After a brief overview of the interclient communications package in this section, its various applications will be illustrated in more detail in subsequent sections, where the basic design and operation of the analysis interface is described.

Some of the information files are accessed by multiple asynchronous processes. One example of a shared information file is the one that stores the values of the interface environment variables discussed above. A token passing scheme is used to ensure that no more that one process has access to any shared-information file at any time. This is more conservative than necessary, as separate tokens could be used for separate files and file access could be specified as either read or write.

Figure 7 illustrates manipulation of the token. Before a new process accesses a shared file it checks to determine if the token is available. The token is an empty UNIX file in this implementation so that the process checks for the existence of the file. If the token exists the process claims it by renaming it to a name based on the process's unique UNIX process identification number. The renaming convention is well known, so that a claimed token can be found with a wildcard file name search.

If a claimed token exists and the process that claimed it exists, the new process will wait and begin the whole procedure of attempting to claim the token again after a brief sleep. If the process that claimed the token has died, an error has occurred. The process that claimed the token has died without releasing it. In this case the process needing a token steals the token although it was never released. This prevents lockups when processes abort unexpectedly before releasing the token. In any case, after successfully claiming the token, the file is accessed and then the token is released.

6.3.3 Launching a Client

As previously stated, all analysis interface commands are stand-alone UNIX executables. A command is executed by launching the appropriate analysis interface client. Appropriately written clients may be launched in one of two modes: *dynamic* or *static*. Any UNIX executable can be configured properly for static launching, but a client must be written especially to support dynamic operation. In dynamic mode, the information displayed by the client changes in response to user interaction with other clients through the medium of the interclient communications package. In static mode the information displayed at start up remains on display until the client is killed, regardless of user actions with other clients. Thus the dynamic version of a client is best used while

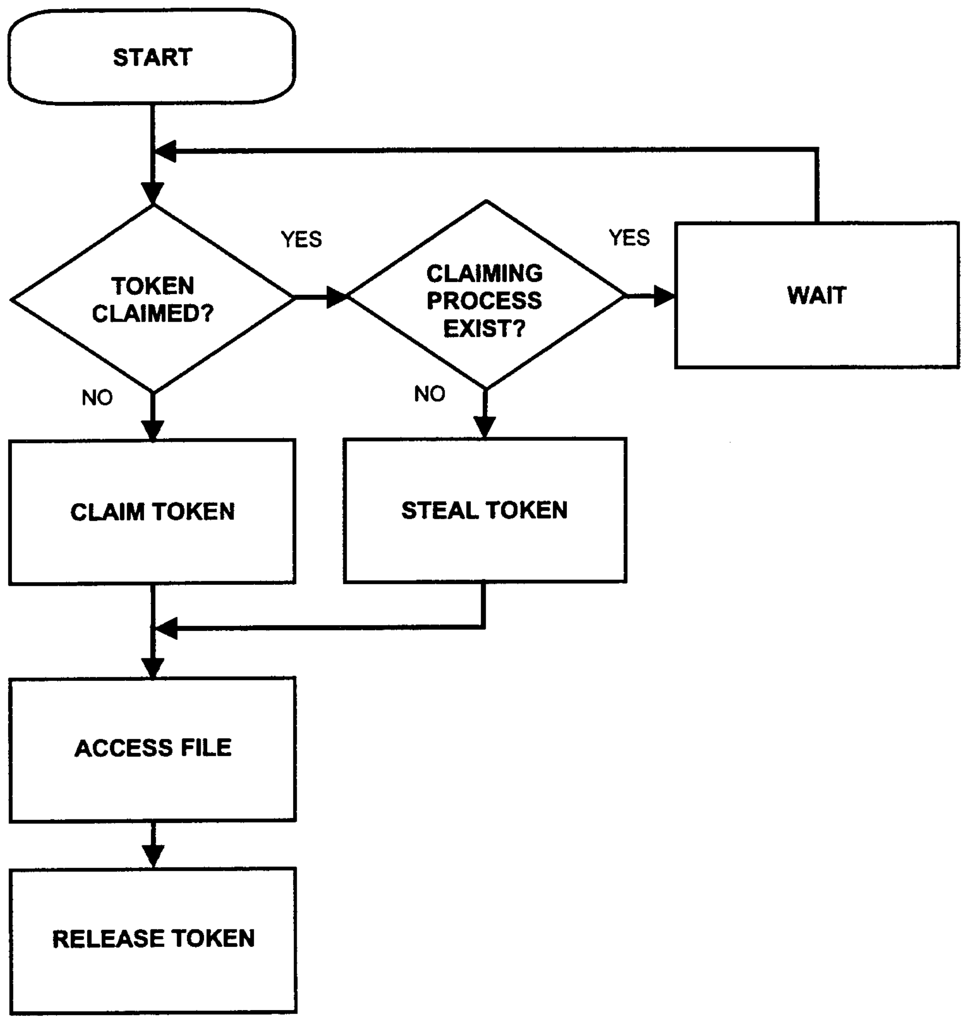

Figure 7. Access to shared information files is synchronized with a token passing scheme.

interactively searching the contents of the databases and the static version of the client is used for intermediate storage of retrieved data.

A client is launched by executing its command file with the appropriate calling sequence. The series of events performed by the command-file process is illustrated in Figure 8. The calling sequence includes the data file names. If the data files do not exist or are out-of-date, then current

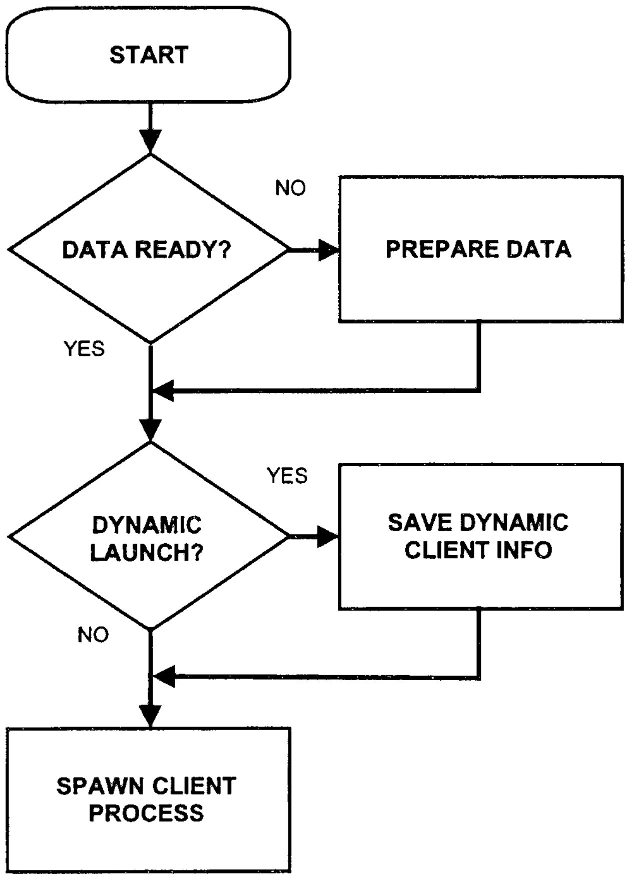

Figure 8. Commands are executed by launching an analysis interface client in either static or dynamic mode.

data files are prepared. Typically UNIX *make* is used to prepare the data files. If the launch is to be dynamic (also determined by the command file calling sequence) then the dynamic nature of the client is recorded by defining an interface environment variable based on the command file process's UNIX process identification number as described in Section 6.3.4. The variable naming convention is well known so that all interface environment variables describing dynamically launched clients can be recovered by a wildcard search of all interface environment variables. This variable will be needed to update the client when the state of the analysis interface changes as described in Section 6.3.5. Statically launched clients will not have this variable defined. The variable value essentially records the calling sequence used to launch the client. The command file process then spawns the client, which inherits the process identification number of the command file process.

6.3.4 Updating the Interface Environment Variables

Typically when a user selects data using one of the analysis interface clients, the client changes the state of the interface by defining the values of one or more interface environment variables. It has also been mentioned previously in Section 6.3.3 that an interface environment variable is defined for each dynamically launched client. To define the value of one or more interface environment variables, the client passes the strings that are to be recorded in the interface environment variables file to a routine that handles the update. The sequence of events performed by the update routine is illustrated in Figure 9.

First, the file is edited to save the new variable values. If the changes will not affect the processing done by any client, then the update procedure is complete. This is the case when a variable is defined that describes dynamically launched clients instead of a new data preference on the part of the user. Otherwise the changes are propagated to dynamic clients. The choice to propagate or not is determined by the calling sequence for the update routine.

In order to propagate the changes to dynamic clients, the update routine retrieves all of the interface environment variables that describe dynamically launched clients and invokes each client's command file with the appropriate calling sequence from information encoded in the variable value. Thus only dynamically launched clients can be updated because only these processes have the appropriate interface environment variable defined. The command file actually handles the updates through a process described in Section 6.3.5.

6.3.5 Updating a Client

A client is updated by executing the same command file that is used to launch the client; however, the calling sequence causes the events illustrated in Figure 10. First, the command file process executes the client's application file that retrieves the interface environment variables, defines a new instance definition for the client, and saves it in the client's instance definition file.

An instance definition is the information contained in a client-dependent text file that describes, among other things, the calling sequence for the client based on the current values of the interface environment variables. The process that defines the instance definition file first reads all of the interface environment variable values, substitutes them in a sequence of command lines and then executes them, writing the instance definition file. The name of the instance definition file is well known and can be read by the client.

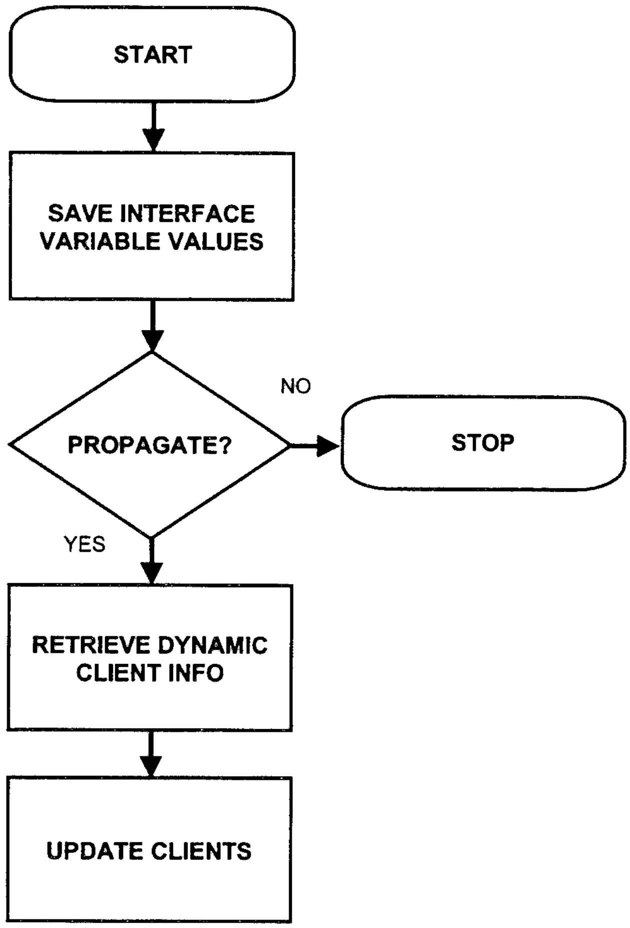

Figure 9. When user interaction or program intervention changes the state of the analysis interface all dynamic clients are considered as candidates for updating.

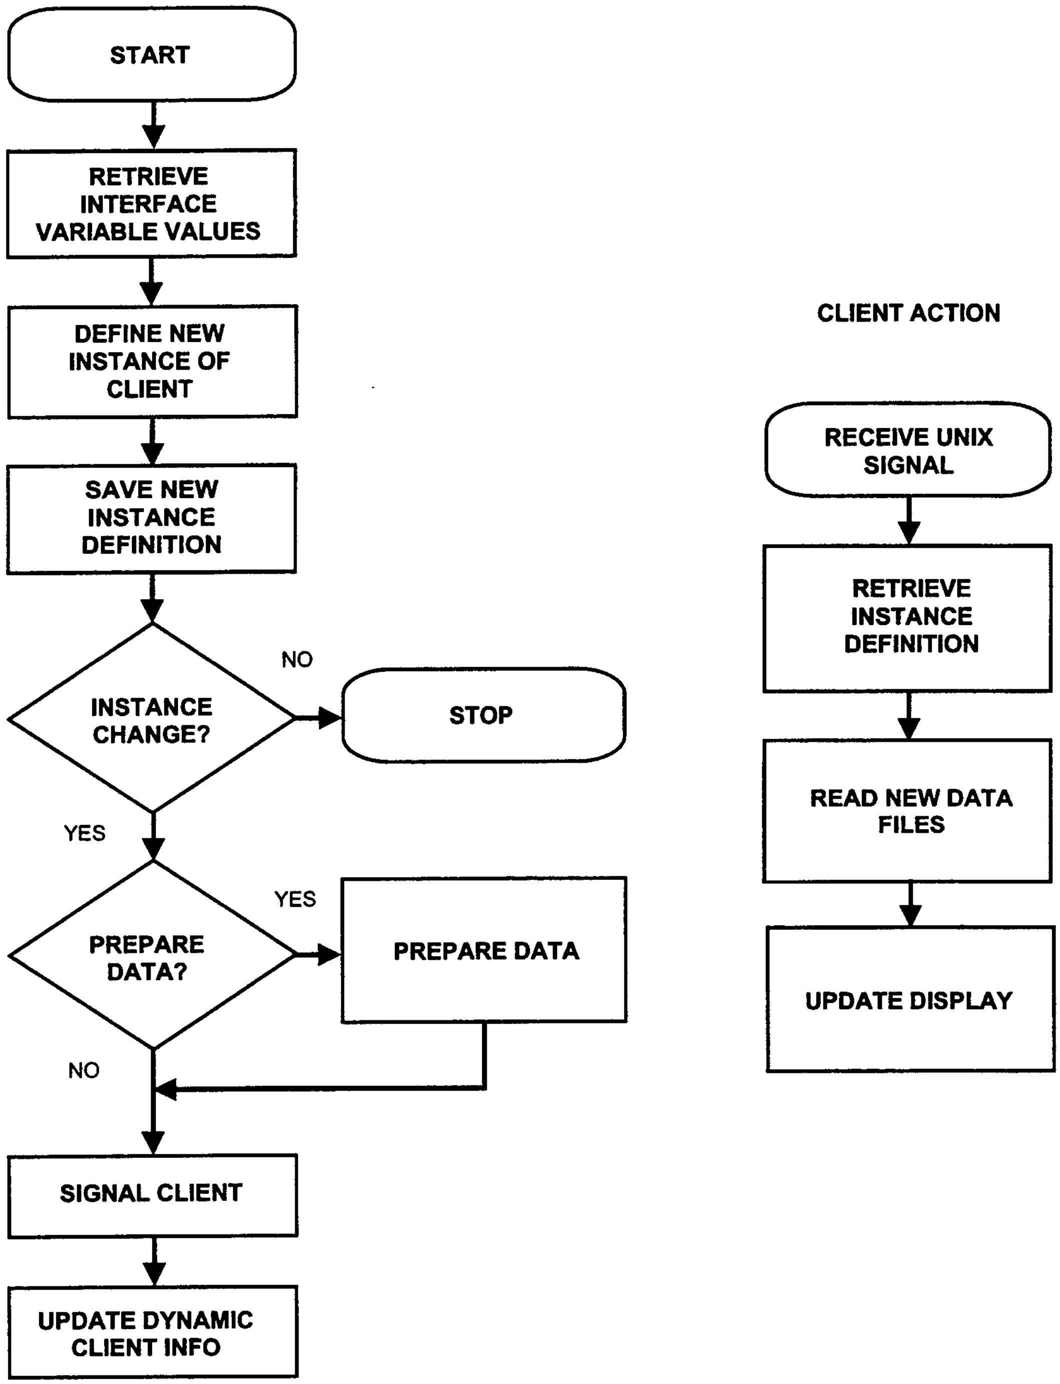

Figure 10. As the state of the analysis interface changes through user interaction or program intervention the dynamic clients update their displays to reflect new user choices.

The old calling sequence for the client extracted from the appropriate interface environment variable is compared to the new one in the instance definition file. If they are the same, then the change in the state of the interface, as described by changes in the interface environment variables, does not affect the client under consideration and the client update procedure terminates. If the calling sequences are different, then the client is not performing the correct processing for the current interface state. In this case any new or out-of-date data files that the client needs based on the new calling sequence are prepared. Typically UNIX *make* is used to prepare the data files. Then the client is signaled with a standard UNIX signal and the new calling sequence is encoded in the client's interface environment variable as described in Section 6.3.4, replacing the old sequence.

When the client receives the signal it retrieves its new calling sequence from its instance definition file and behaves almost as if it was executed with the new calling sequence. This could be trivially done by spawning a new process, but by using that simple approach the history of the user's interaction with the client and its window manager up to the time of the update would be lost. Instead, the client typically initializes certain internal options, leaves others as reset by the user, and reads and displays new data files.

6.4 Adding Clients to the System

It should be clear from the previous discussions that much of the overhead of providing interaction between the clients of the analysis interface is distributed among the clients themselves. Because of this distributed responsibility, new clients may also be added to the analysis interface without changing its basic structure. One simply needs to provide the new client with the command and application files. There are basically three scenarios for adding new clients. Standard X Windows and UNIX programs may be incorporated so that they can be launched by the *CommandUtility* (in static mode only). Clients based on existing dynamic clients (such as *xgraph*) may be added. This would, for example, allow the user to add a new graph displaying different information than the existing graphs. Entirely new clients may also be written with dynamic update capability and incorporated into the interface. In order for a client to be added, the application file and the command file must be provided. Of course, these must be properly implemented with the synchronization protocol in place and with proper utilization of the interface environment variables. Instructions and examples of adding clients are included with the software distribution.

7. CONCEPT OF OPERATIONS

Once the facility is properly set up, it operates autonomously and does not require human intervention to process signature data or to maintain its databases. The operational concept for satellite anomaly detection, however, requires a human analyst in the loop to subtly alter the input data to the system and to notify users of confirmed satellite anomalies. The analyst typically becomes involved with a few percent of the observations processed by the system. The observations of interest to the analyst are the ones flagged by the system. The external manifestation of a flag is an electronic mail message that the system mails to analysts on its distribution list. Figure 11 shows how the sensor, the system, and the analyst interact to provide satellite anomaly detection information from the SSN satellite signatures.

The process, which is illustrated in Figure 11, begins with the tasking supplied to the SSN. A large number of satellites must be tracked on a regular basis in order to maintain the catalog of satellite orbits. Each single track of a satellite is referred to as an observation. The signatures collected during catalog maintenance of HEO satellites should all be sent to the processing system. The system is typically configured in such a way that a few percent of the observations are flagged, based on the amount of cross-section fluctuation observed in the signature and the history of the satellite. There are two types of flags. A flag will be either a rare signature classification event or a satellite characterization change detection. The type of flag and other relevant information are sent as email to analysts on a distribution list maintained on a system file.

The analyst should attempt to determine, on a flag-by-flag basis, if an anomalous satellite behavior has occurred. The first step in this process probably requires that the analyst visually examine the signature associated with the flagged observation. If it is a variable cross-section signature, could the cross-section fluctuation be caused by sensor-tracking errors or interference from some external noise source? Is the fluctuation typical of slow satellite aspect angle change or could the satellite be tumbling? An experienced analyst will be able to identify many flags that he or she can explain by some phenomenon other than anomalous satellite behavior.

Of the observations that remain unexplained except by anomalous satellite behavior, some will be associated with satellites that have appeared to exhibit the anomalous behavior over several previous observations, and the analyst will conclude that indeed the status of the satellite has changed. If the anomalous behavior has been evidenced in only one observation, the analyst might hypothesize that the status of the satellite has changed, but he or she will want to request another one or two observations of the satellite to confirm the hypothesis. This will occur for only a fraction of the observations flagged by the system, probably only a few times per one-thousand satellite observations.

The interaction between the analyst and the system can be illustrated with an example provided by the accidental destabilization of a payload. Figure 12 shows the expected probability of observing a constant cross-section signature for the payload from December 1993 through September 1994 as computed from the narrowband signatures of the Millstone Hill L-band radar. Each symbol

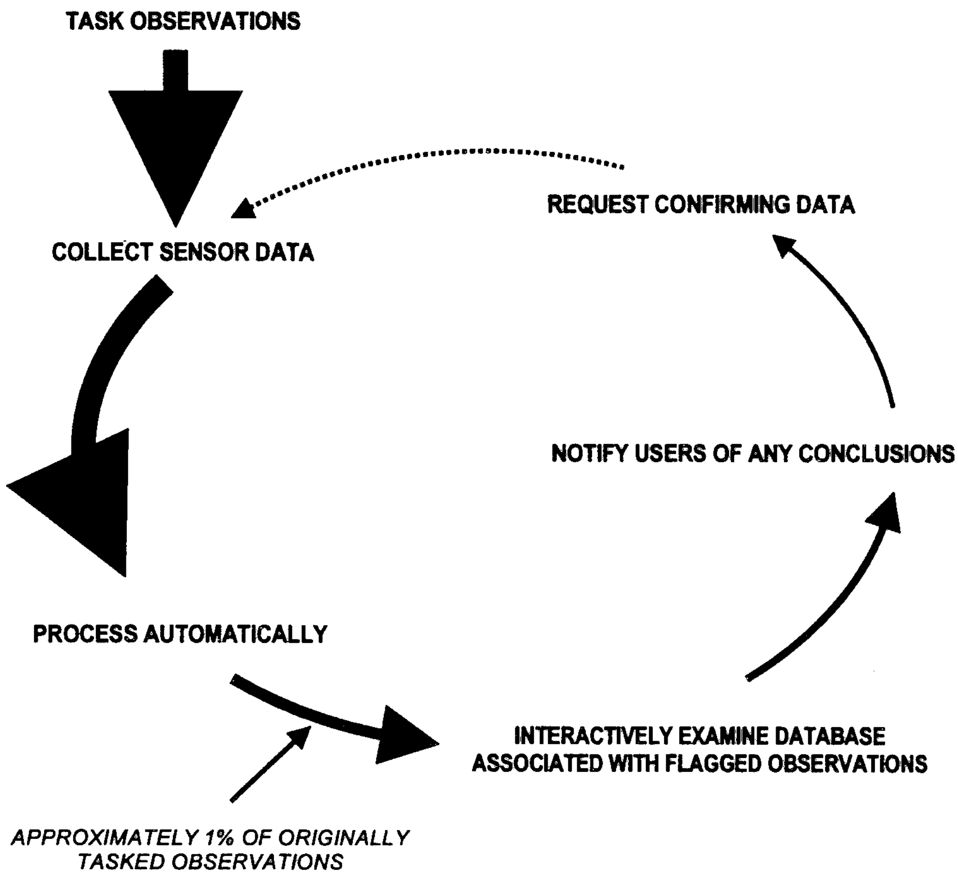

Figure 11. The operations concept includes an automated signal and data processing system that processes all available satellite signatures and notifies a human analyst of anomalous satellite observations (typically about 1% of all observations). The analyst may interactively examine the database to determine if the anomaly is of interest and either notifies users or requests more data to verify his or her hypotheses. The line weights used in the arrows that depict information flow are approximately proportional to the data volume at that point.

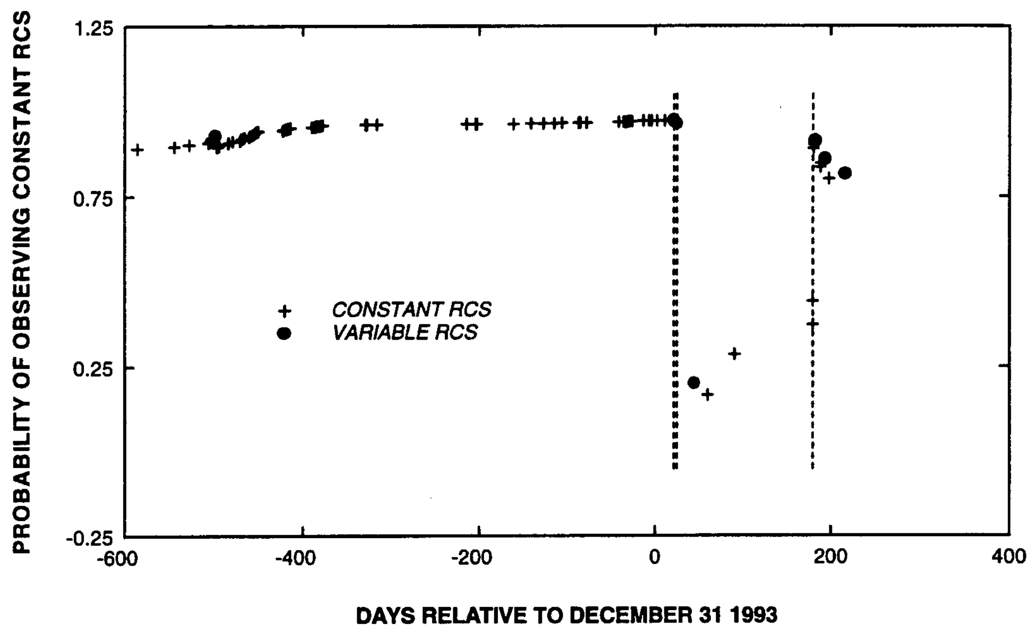

Figure 12. Three anomalous observations (marked by the three vertical lines) provide the analyst with the starting points that finally lead to recognition of the destabilization and subsequent restabilization of a satellite payload.

represents a single observation of the satellite. The three vertical lines pass through the three observations that were flagged by the system as anomalous. The first two lines from the left are close together, separated by only three days.

The first vertical line (from left to right) marks an observation made on day 21 of 1994. The signature was classified as a rare, variable cross-section signature. In the system's message to analysts it was noted that the previous 73 observations were associated with constant cross-section signatures. If a possible change in the satellite's stability was of interest, the analyst could have examined the flagged signature shown in Figure 13(a). The cross-section fluctuations reach a peak sixteen times the average value observed. This could be due to rotation of the satellite if it is not an artifact of the sensor observation. The analyst could ask the sensor to retrack the satellite for confirmation.

The sensor did track the satellite again three days later. This observation is marked by the second vertical line from the left in Figure 12. The system again flagged the observation as exhibiting a rare, variable cross-section signature. The probability of two variable cross-section signatures in a row was so small, considering the history of the satellite, that the observation was also flagged as a satellite characterization change detection. The second variable cross-section signature is shown in Figure 13(b). This signature shows features similar to those seen in Figure 13(a).

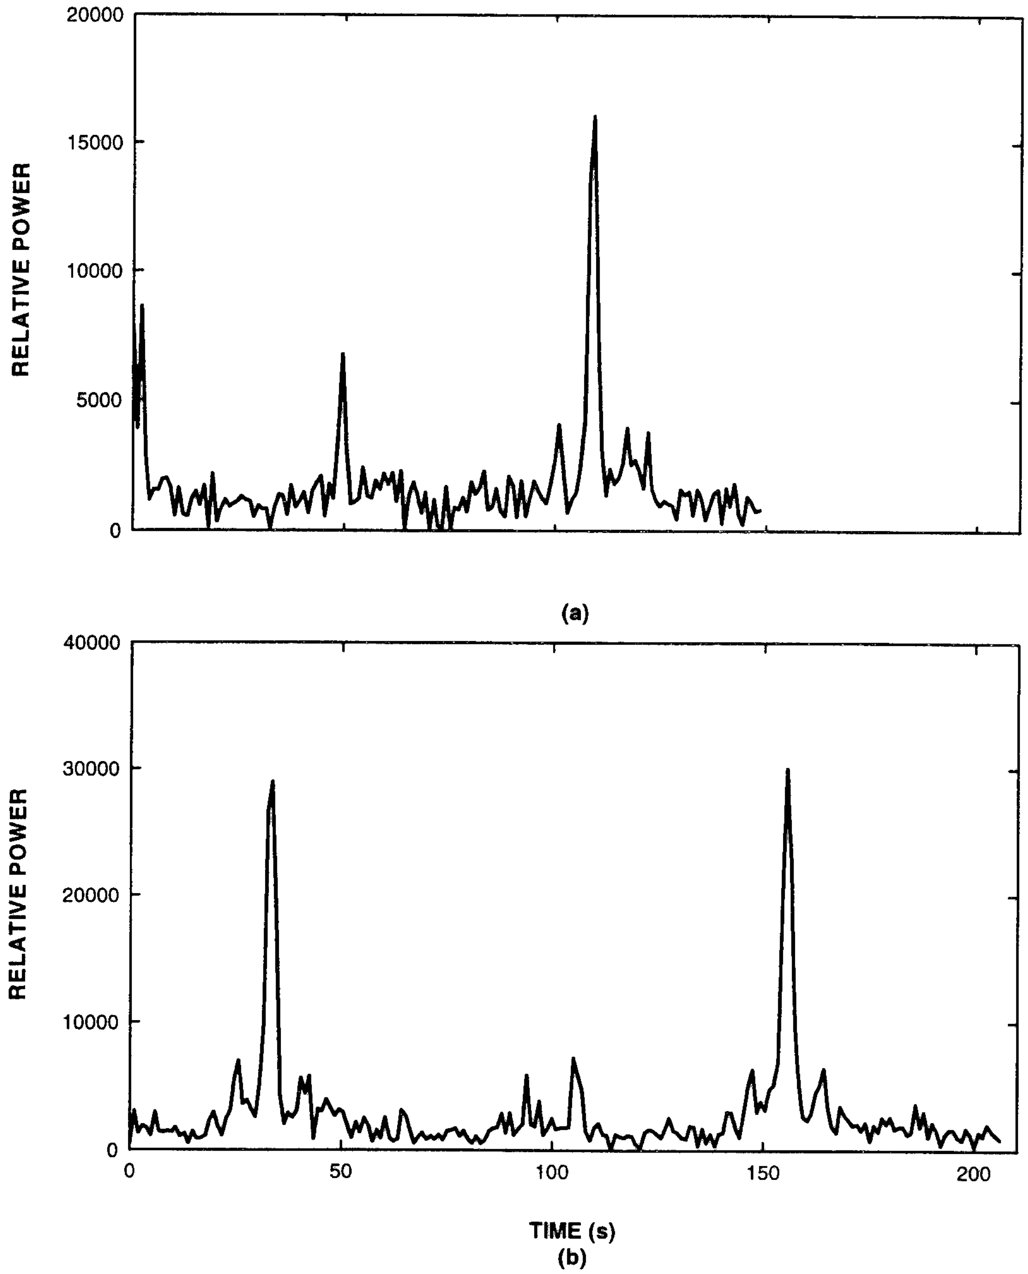

Figure 13. The analyst uses the analysis interface to interactively examine satellite observations flagged by the system's anomaly detection algorithms. In (a) the first rare, variable cross-section signature flagged by the system is shown. The previous 73 signatures were of the constant cross-section type. This is the first indication to the analyst that the satellite stability pattern has changed. In (b) the second rare, variable cross-section signature flagged by the system is shown. This signature also caused a satellite characterization change detection flag because, given the satellite's history, two variable cross-section signatures in a row is a very small probability event.

In fact, the satellite had transitioned from a three-axis stabilized condition to a tumble, and the system easily detected the change. If there had not been enough data to establish that the satellite was normally three-axis stabilized, or if the change in the status of the satellite had not caused it to tumble, then the observations would not have been automatically flagged. That is, the change detection algorithm provides an automatic indication of a change in the satellite's stability pattern.

The satellite was successfully restored to its three-axis stabilized condition sometime after the middle of February. The system needed to observe four constant cross-section signatures in a row in order to declare the characterization change detection because the satellite was not tumbling long enough to precisely characterize the new probability of observing a constant cross-section signature after stabilization was lost.

8. SYSTEM PERFORMANCE

This section describes some of the characteristics of the satellite signature database collected from the Millstone Hill L-band satellite tracking radar and processed by the software described in this report. The database characteristics were compiled shortly before publication of this report. The radar data were collected over a period of several years, excluding only the small fixed database used in the initial pilot study for the project. The earliest processing was done by the engineering prototype for the system and the database was moved to the release-ready software sometime in May of 1994.

First, the various subsets of data are described. Figure 14 provides some insight into the size of the database, which includes 49,993 observations of 1,482 satellites, broken down into various subsets. The largest set of data includes all observations sent to the system. Not all of these observations warrant further processing. The system employs a simple method for rejecting bad data in order to reduce the chances of corrupting the database. Currently there are two tests.

The first test checks the consistency between the orbital element set used to compute the values of various geometric variables during data processing and the observed sensor state vector. Specifically the expected direction from the sensor to the satellite (called the pointing) is computed from the orbital elements and compared to the analogous quantity estimated by the sensor's target tracking software during the satellite observation process. Data may be rejected if the differences between these predicted and observed pointing quantities become large enough. The second test applies the same randomness test used for signature classification to the noise power-time series supplied by the sensor as part of the observation. At the Millstone sensor the two tested noise power sequences correspond to the power that would be returned by a target just slightly closer to the sensor, and slightly farther away from the sensor, than the real satellite is estimated to be. If the target is really where the sensor thinks it is, then these noise sequences will be (stationary) random processes. If the sequences appear nonrandom, or fluctuating, then some target signal may be affecting the noise sequences and the target may not be at the range estimated by the sensor. This tracking error could cause the satellite signature to fluctuate even though the satellite cross section is really constant, so such data is rejected.

The purpose of the first test is to avoid using orbital element sets that do not go with the target. The purpose of the second test is to reject signatures that fluctuate due to tracking error instead of due to an actual characteristic of the satellite cross section. The satellite observations that are not rejected by the previous two tests are called consistent observations—the noise channel statistics, satellite element set, and the state vector reported by the sensor's tracking system basically agree. Figure 14 indicates that 43,581, or about 87%, of the original observations have passed both of these tests. This is not really a large rejection rate for a sensor that has been calibrated primarily for making the measurements required to deduce satellite orbits instead of measuring satellite signatures. Because the system was designed to accept all of a sensor's observational data,

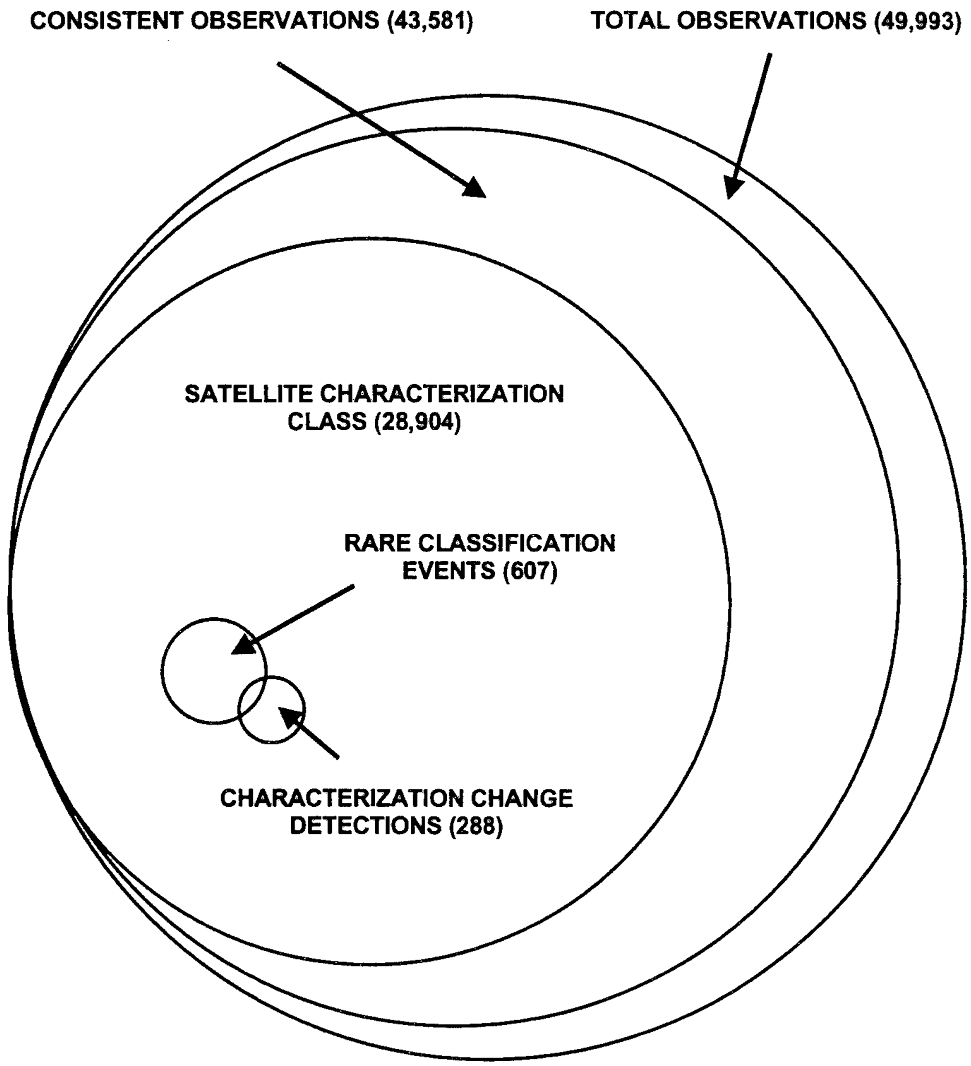

Figure 14. The different groupings of satellite observations are shown for a typical database. The area of each circle is proportional to the number of observations in the portion of the database that it represents. The numbers in parentheses are the actual number of observations.

it is necessary from a practical viewpoint to empower the system with some data rejection capability. Although no further effort is made to find suspicious or poor quality observations, not all of the remaining data can be used to characterize the satellite. The next reduction in the quantity of data is designed to both avoid the problem discussed in Section 4.4 and raise the probability of detecting anomalous observations with the algorithms described in Section 5.2. Depending on the satellite, the data must satisfy certain observation conditions in order to be further processed for the purpose of characterizing the satellites. Observations satisfying one of these sets of conditions are said to be members of a satellite characterization class.

Some research was done on automated definition of the appropriate observation conditions for each class in Raup [5], pages 23–30, but the appropriate algorithms were never fully developed and implemented in the system. Currently the observation conditions are statically defined in software and define two characterization classes. A given satellite is associated with one of the characterization classes based solely on its orbit. A typical scheme for determining the characterization classes is shown in Figure 15.

Both characterization classes reject observations with phase angle values near 0 and 180 deg. The phase angle is defined on page 11 of this report. When the phase angle is nearly 0 or 180 deg then the sun, the satellite, and the sensor are almost located along a line, so this condition is referred to as solar alignment. The exact thresholds determining solar alignment can be configured. Because of the sensitivity of a satellite's cross section to small aspect angle changes when viewing the satellite along lines nearly perpendicular to the plane of large solar panel arrays, signature fluctuation is often observed near solar alignment. This observation is often made even if the satellite is inertially stabilized and shows no evidence of variable cross-section signatures at any other phase angles. The condition occurs in properly functioning satellites that maintain the alignment of solar panels to maximize power production.

Figure 3 shows both satellites exhibiting variable cross-section signatures near solar alignment even though one of the satellites is thought to be three-axis stabilized and one satellite is thought to be tumbling. Therefore, cross-section fluctuation measurements made near solar alignment would be less effective in distinguishing between the two stability states of the satellites than fluctuation measurements made at other phase angles. By eliminating the observations collected near solar alignment, the change detection algorithm can more effectively detect when the behavior of a satellite changes from one of the illustrated stability states to the other. As discussed in Section 4.2, the nonrandomness of the fluctuation measurements for the satellite in Figure 3(b) is eliminated by restricting the observations so that they are not collected near solar alignment.

One characterization class uses all of the consistent observations that are not collected near solar alignment. All HEO observations of satellites in nearly circularly shaped orbits, as opposed to very elliptically shaped orbits, are tested for inclusion in this class. However, if this were the only condition satisfied by the observations from satellites in very elliptical orbits, then nonrandomness due to the effect discussed in Section 4.4 would not be eliminated. The signal-to-noise ratios associated with the observations would vary a large amount, depending on the orbital position.

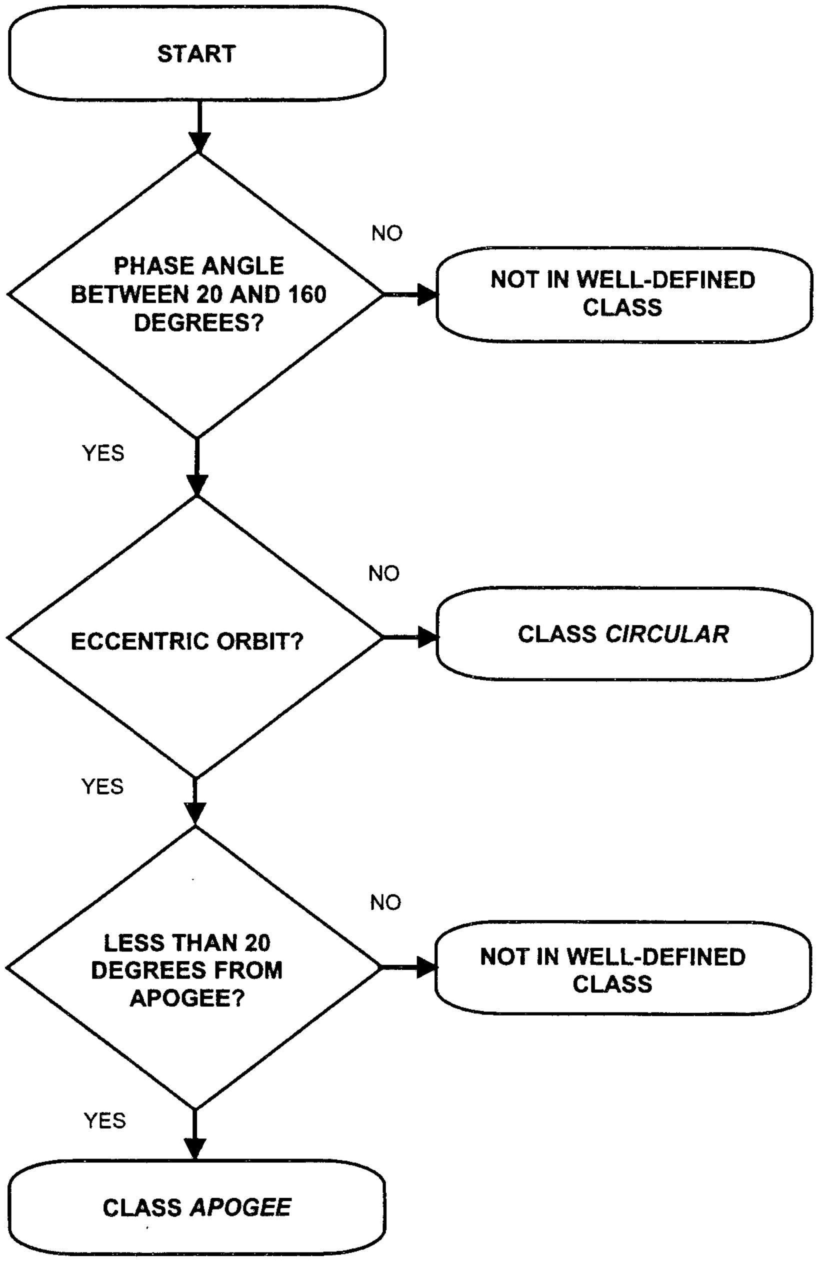

Figure 15. The characterization classes contain observations of satellites satisfying certain conditions that increase the sensitivity of the anomaly detection algorithms.

For this reason a second characterization class is defined for satellites having very elliptical orbits. As with solar alignment, the thresholds for whether a satellite is considered to be in a circular or elliptical orbit is site dependent.

The second characterization class, which applies only to satellites in highly elliptical orbits, uses all of the consistent observations that are not collected near solar alignment but are collected near the apogee of the orbit. The decreased signal-to-noise ratios associated with the signatures of satellites near apogee make it more difficult to detect variable cross-section signatures, but satellites placed in such orbits often exhibit variable cross-section signatures away from apogee regardless of whether they are stable at apogee or not. This can be seen in Figure 16(b). Only within a few degrees of apogee (180 deg of true anomaly) does the satellite consistently exhibit constant cross-section signatures even though it is thought to be successfully three-axis stabilized. The data in Figure 16(a) is from a similar satellite that is no longer three-axis stabilized. By eliminating the observations collected away from apogee, the change detection algorithm can more effectively detect when the behavior of a satellite changes from one of the illustrated states to the other. As discussed in Section 4.2, the nonrandomness of the fluctuation measurements for the satellite in Figure 16(b) is eliminated by restricting the observations so that they are not collected away from apogee.

A signature must be a member of at least one of the characterization classes in order to contribute to the characterization of a satellite or to be flagged by one of the two anomaly detection algorithms as described in Section 5.2. Besides the signatures that are flagged by the anomaly detection algorithms, another subset of the signatures belonging to a characterization class must be considered. The characterization change detection algorithm is designed to determine when the underlying probability of observing a constant cross section has changed within an observation class of signatures. Hence, a positive characterization change detection result is used to discard old observations before the declared characterization change and initialize a new satellite characterization attempt. All observations before the event that triggered the positive change detection are discarded for purposes of characterizing the satellite. This is how the system records changes in the behavior of the satellite as the satellite progresses from birth to death. When such discarded observations are removed from a characterization class of observations, the remaining observations are said to be in the current data set. All current data sets, except the first one (when the current data set and the characterization class of observations are exactly the same set of signatures), contain exactly one observation flagged by the characterization change algorithm.

The anomaly detection algorithms operate with a small fixed probability of false alarm. Therefore nearly all positive anomaly detections are real. However, an anomalous observation need not be related to a significant change in the status of the satellite. The purpose of the anomaly detection algorithms is to reduce the number of signatures that need to be reviewed by a human analyst, as explained in Section 7, which describes the concept of operation, including the roles of the sensor, the system, and the analyst. Figure 14 shows that only about one percent of signatures are flagged as satellite characterization change detections. The other 99% of signatures are usually of little interest to the analyst. The actual probabilities of space surveillance event detection and

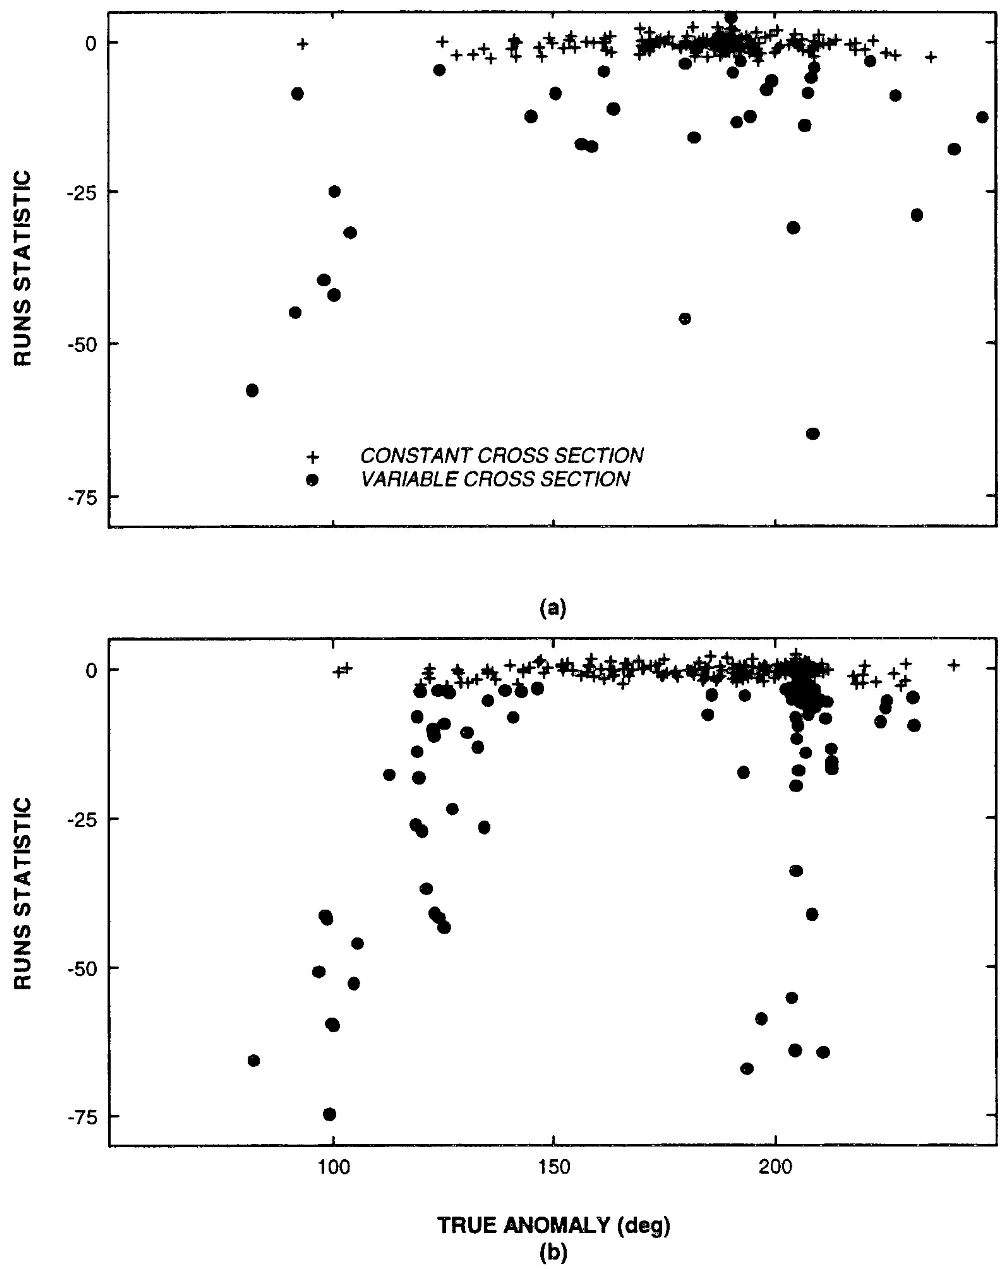

Figure 16. When a satellite operates in a highly elliptical orbit, only signatures near apogee are used to characterize the satellite. In (a) it can be seen that a satellite that is not three-axis stabilized exhibits variable cross-section signatures at all indicated values of true anomaly. In (b) it can be seen that the satellite, which is three-axis stabilized, exhibits variable cross-section signatures only away from apogee (180 deg of true anomaly). The change detection algorithm is made more sensitive to change from a three-axis stabilized state to a tumbling state if only the signatures near apogee are used to characterize the satellite.

miss clearly depend on what events the user is considering. The system is best at characterizing the stability history of a satellite and flagging changes in stability on a satellite-by-satellite basis. The rate at which observations are flagged is small enough to apply a real bulk-filtering effect to the large number of signatures produced by the space-surveillance system. In the next part of the section the ability of the system to find satellite characterization change detections and rare signature classification events is investigated.

A change detection flag is raised when a particular signature classification type, either constant or variable cross section, is observed repeatedly and the number of repetitions is of sufficiently small probability with a given confidence. How many signatures of the same type must occur in a row in order to raise the change detection flag? This depends, of course, on the observation history of the satellite. Figure 17 shows the number of satellites requiring N or fewer observations with signatures of the same classification type in order to trip the change detection alarm. It is a measure of the readiness of the system to detect extreme changes in satellite stability. The data used to produce the graph has been restricted to satellites with at least one observation in its current data set. When a satellite is three-axis stable then most of the signatures from its satellite observation class will generally be classified as constant cross section signatures. So if the satellite is flagged by the change detection algorithm it is because one or more variable cross-section signatures were consecutively observed.

Notice that more than 400 satellites would be flagged by a run of four or less signatures of the appropriate classification. Figure 14 has already shown that only about one percent of signatures in a satellite's observation class lead to a change detection. Therefore, the change detection monitoring process works fairly well. The system is ready to flag satellites that change their stability patterns, but most satellites maintain the same stability pattern over their normal operational lifetime.

The rare classification event test automatically flags signatures of a particular type, either constant or variable cross section, that occur less than a specific fraction of the time with a given confidence. It may be that a satellite does not have a particularly large or small probability of observing a constant cross-section signature, so that neither type of signature is statistically rare. This situation depends only on the history of the satellite. If one type of signature would be flagged by the system as rare, then the rare classification event test is said to be *armed* for that satellite. Figure 18 shows the number of satellites with armed rare classification event tests.

This section ends with a brief profile of the raw observation database in order to provide an indication of the typical observation characteristics of a useful observation data set. The purpose of Figure 19 is to present some idea of the characteristics that have led to the performance described in this section. Two of the most important characteristics of the signature database are the number of signature samples and their signal-to-noise ratios. The number of signature samples directly determines the population size in the signature classification tests. The signal-to-noise ratios of each sample determine how noise-like the signature will appear when there is true cross-section variation. Because the signature classification algorithm operates at a fixed probability of erroneously declaring a constant cross section signature to be variable, both attributes determine the ability of the system to correctly identify the variable cross-section signatures.

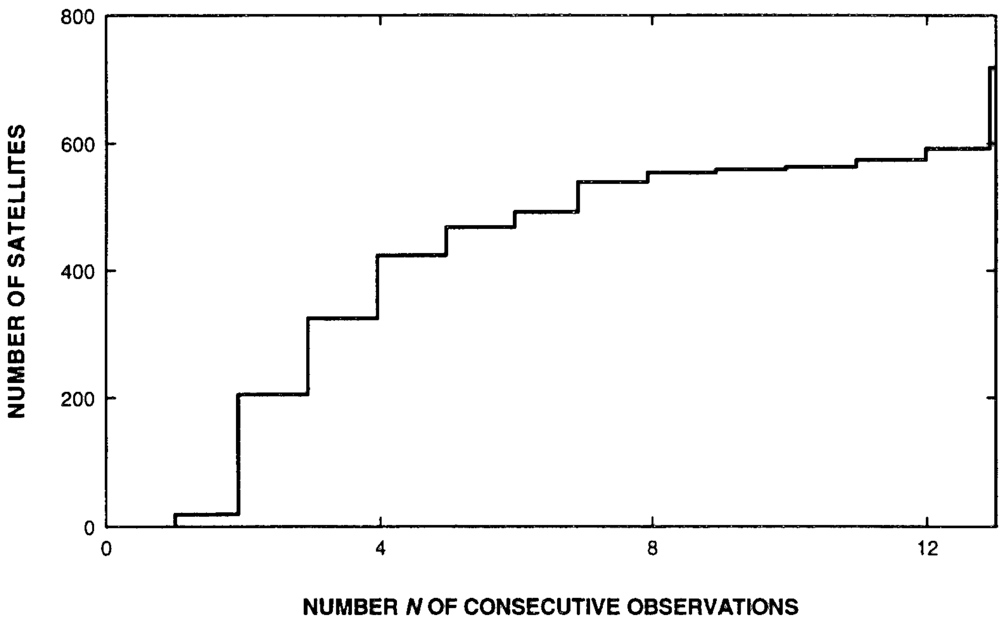

Figure 17. The number of satellites requiring N or fewer consecutive signatures of the appropriate classification type to trigger a characterization change detection flag is shown. More than 400 satellites would be automatically flagged by a run of four or fewer signatures of the same appropriate classification type. Thus these satellites are being efficiently monitored for extreme changes in their stability patterns.

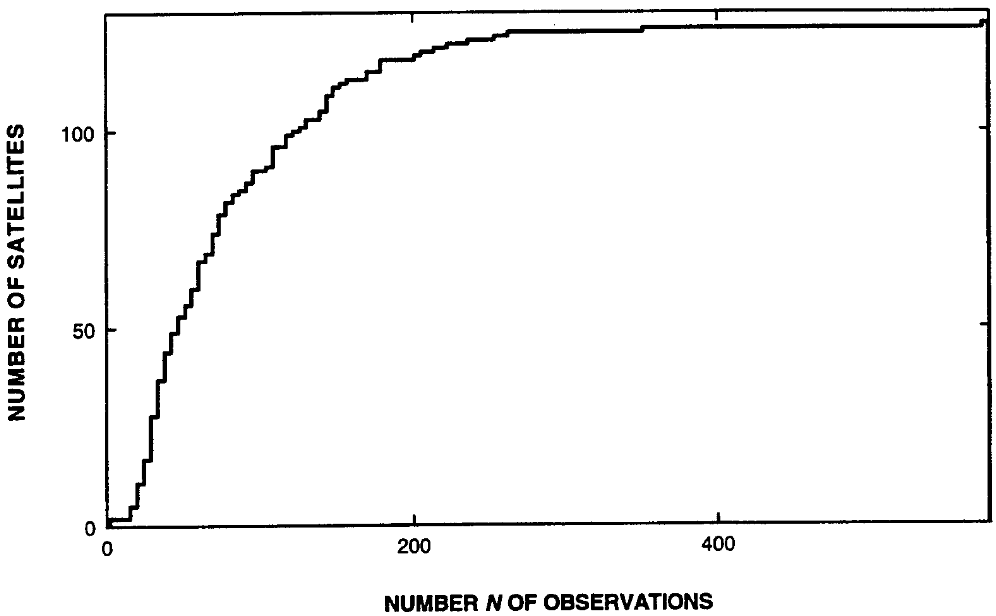

Figure 18. The number of satellites with armed rare event tests having N or fewer observations in the current data set are shown. It can be seen that most satellites with armed tests have 75 or fewer observations in their current data set. This indicates the size of the current data set required to perform rare signature classification event detection.

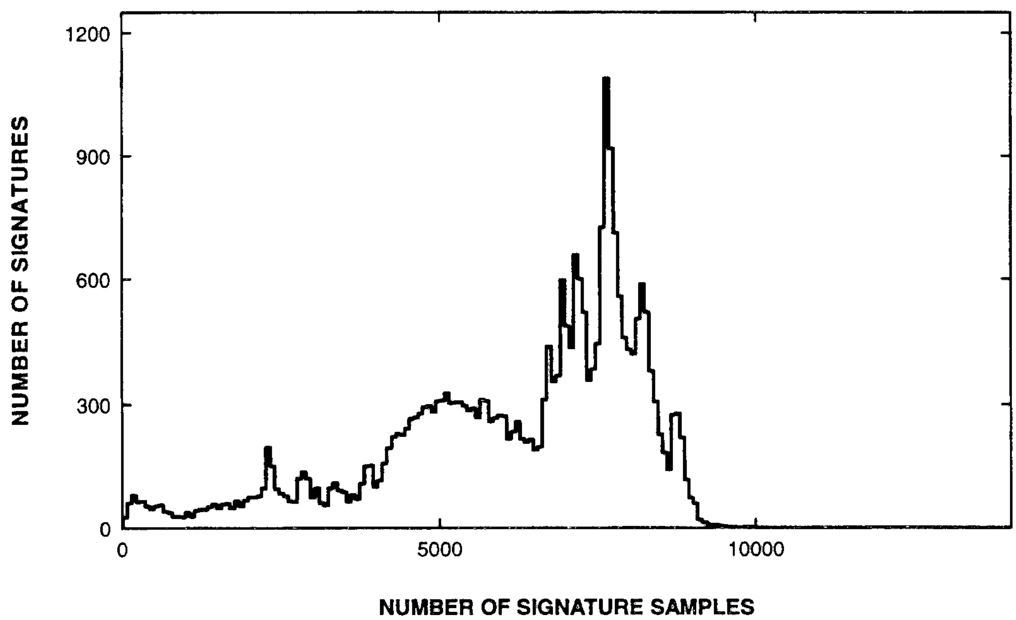

(a)

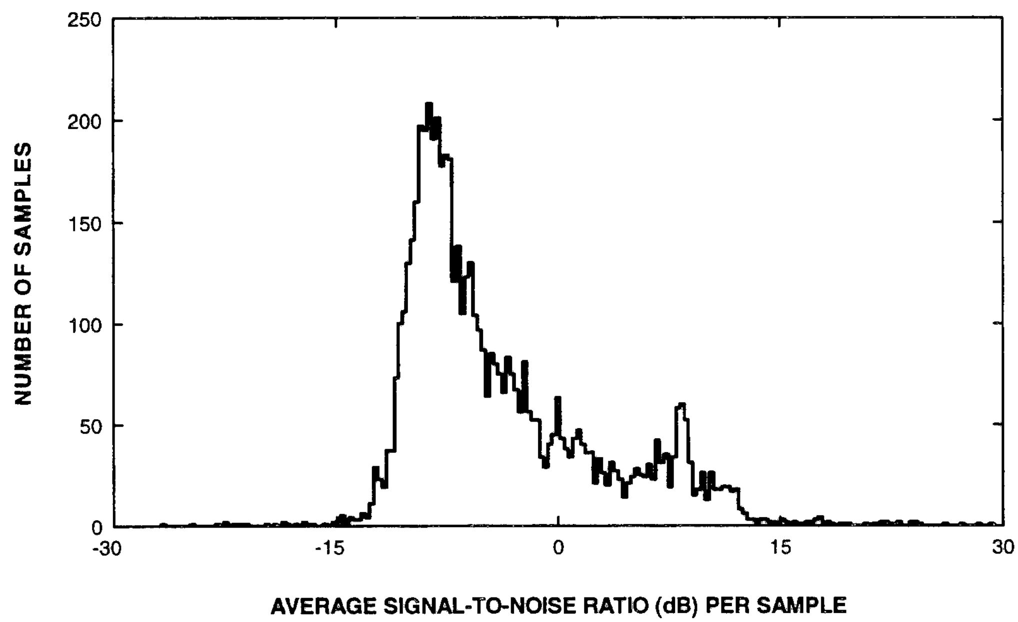

(b)

Figure 19. Two histograms (density estimates) describing important characteristics of the raw signature observations collected from the test sensor are shown. In (a) it can be seen that several thousands of signature samples are collected in most satellite observations, representing several minutes of data collection for each observation. In (b) it is seen that the samples are generally taken at negative signal-to-noise ratios.

9. SUMMARY

Measurements from the sensors in the SSN are used to maintain a catalog of the orbits of thousands of satellites. Because of the large amount of data, techniques have been developed over the years to automate significant portions of the catalog maintenance process. As the position of each satellite is measured by a radar or optical sensor, the sensor also produces a satellite signature. A signature is a time series that describes the microwave or optical power, depending on the sensor, that is reflected from the satellite as a function of time. These signatures contain valuable information about the satellite cross section but they are not routinely used.

Now it is possible, at least in the regime of HEO, to automatically process satellite signatures from the SSN and significantly improve the ability of the network to monitor the operational status of satellites. The techniques are suitable for use with the signatures already available from the routine catalog maintenance operations, in spite of some known problems with sensor measurements. Experience shows that it is difficult to maintain the amplitude calibration of a sensor over a long period of time, or even to determine the state of a sensor's amplitude calibration by examination of its data. For this reason the system measures the amount of cross-section fluctuation rather than the absolute cross-section levels. This approach allows amplitude calibration biases to vary from track to track without degrading system performance, providing robustness across sensors in the network from day to day.

A UNIX-based software system has been developed, built, and operated at the Millstone Hill L-band satellite surveillance radar, a contributing SSN sensor. Nearly 50,000 signatures from HEO satellites have been processed by the system, which is now ready for distribution to other space surveillance systems and data analysis centers. Only generously licensed software is used in the system, making distribution and maintenance of the software more economical. The developer's system will continue to operate at the Millstone Hill sensor. Its modern network-based design allows signal and data processing components and its graphical user interface to be located at geographically diverse sensor and analysis sites as appropriate. Access to the system is even routinely achieved with a desktop PC-compatible computer, a modem, and a conventional phone line. Response time is slow in this configuration because of the small bandwidth of the conventional phone line, but future improvements in communications equipment and software, or simply using a wider bandwidth connection today, will alleviate this response-time problem.

The system consists of three components. The signal processing component is primarily used to compress the raw signatures into an archival database. The data processing component uses the signatures to characterize each satellite and then, through a process called anomaly detection, flag unusual observations or changes in the satellite characterization. All of the database processing concepts are based on measurements of the amount of cross-section variation that occurs as the satellite is observed under various conditions. Finally, a graphical user interface component is provided so that human analysts can respond to the anomaly detection messages by interactively examining the data with various tools. The interface also contains a number of novel features that provide flexibility for advanced users without precluding use by novices.

The system has demonstrated that it is capable of bulk filtering the complete signature output of a space surveillance sensor. Even though only about one percent of signatures are flagged for further analysis by human analysts, the system is capable of automatically isolating significant changes in the stability behavior of individual satellites. A three-axis stabilized satellite generally shows very little cross-section variation except under special observation conditions. Without effective three-axis stabilization a satellite's cross section is often observed to vary. The system is designed to flag a change in a satellite's operational status that is accompanied by a significant change in the pattern of its stability behavior.

REFERENCES

1. A.K. Bhattacharyya and D. Sengupta, *Radar Cross Section Analysis and Control*, Norwood, MA: Artech House (1991).
2. N. Levanon, *Radar Principles*, New York, NY: John Wiley and Sons (1988).
3. M.H. Carpentier, *Radars: New Concepts*, London: Gordon and Breach Science Publishers (1968).
4. M. Born and E. Wolf, *Principles of Optics*, Elmsford, NY: Pergamon Press (1983).
5. R. Raup, "Bayesian characterization and detection of rare binary signature events," MIT Lincoln Laboratory, Lexington, MA, Technical Rep. 876 (1990). DTIC AD-A221116.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	7 September 1995	Technical Report

4. TITLE AND SUBTITLE

A Satellite Anomaly Detection System that Uses Signatures from the Space-Surveillance Network of Sensors

5. FUNDING NUMBERS

C — F19628-95-C-0002
PR — 77
PE — 35910F

6. AUTHOR(S)

Richard C. Raup, J. Scott Stuart, and John W. Curtis

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Lincoln Laboratory, MIT
244 Wood Street
Lexington, MA 02173-9108

8. PERFORMING ORGANIZATION REPORT NUMBER

TR-1020

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Electronic Systems Center
Hanscom AFB
Bedford, MA 01731-9108

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

ESC-TR-95-003

11. SUPPLEMENTARY NOTES

None

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

Lincoln Laboratory has developed a capability to assess and monitor the status of high-earth-orbit (HEO) satellites using simple radar-derived satellite target signatures. The capability takes the form of a computer workstation that processes, stores, and retrieves HEO satellite signatures and associated data products. The important sensor phenomena, signal and data processing algorithms, software architecture, and concept of operations are all described in this report.

The system was designed and tested using L-band radar data but will accept other narrowband radar data and photometric data as well. Signal and data processing algorithms are based on statistical modeling and methods. The amount of cross-section fluctuation in each signature is measured as a function of various independent observation variables. The fluctuation measurements are first used to characterize the normal operation of the satellite and then to provide a basis for automatic detection of anomalies. Thus changes in satellite status that manifest changes in the apparent stability of the satellite are detected by the system. The algorithms are embedded in a network-transparent UNIX software architecture with an MIT X Windows and Motif interface. The software for the system can be used on any UNIX platform with little or no modification.

The basic facility may consist of one or more contributing radars or photometric sensors, a data processing center, and one or more analysis centers. These facility elements need not reside at the same location because the software may be distributed among several computing platforms at diverse geographical locations. Signal processing, data processing, data base management, and anomaly detection are automated, while data base retrieval has both an automated mode for sensor control-room displays and an interactive mode for data analysis.

14. SUBJECT TERMS

space surveillance
narrowband radar
photometry
signature analysis
anomaly detection
UNIX software

15. NUMBER OF PAGES

80

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	Same as Report